arXiv:1612.05110v1 [cs.DB] 15 Dec 2016

Efficient Detection of Complex Event Patterns Using Lazy Chain Automata

Ilya Kolchinsky
Technion, Israel Institute of Technology

Assaf Schuster
Technion, Israel Institute of Technology

Danny Keren
Haifa University

Abstract

Complex Event Processing (CEP) is an emerging field with important applications in many areas. CEP systems collect events arriving from input data streams and use them to infer more complex events according to predefined patterns. The Non-deterministic Finite Automaton (NFA) is one of the most popular mechanisms on which such systems are based. During the event detection process, NFAs incrementally extend previously observed partial matches until a full match for the query is found. As a result, each arriving event needs to be processed to determine whether a new partial match is to be initiated or an existing one extended. This method may be highly inefficient when many of the events do not result in output matches.

We present a lazy evaluation mechanism that defers processing of frequent event types and stores them internally upon arrival. Events are then matched in ascending order of frequency, thus minimizing potentially redundant computations. We introduce a Lazy Chain NFA, which utilizes the above principle, and does not depend on the underlying pattern structure. An algorithm for constructing a Lazy Chain NFA for common pattern types is presented, including conjunction, negation and iteration. Finally, we experimentally evaluate our mechanism on real-world stock trading data. The results demonstrate a performance gain of two orders of magnitude over traditional NFA-based approaches, with significantly reduced memory resource requirements.

1 Introduction

The goal of Complex Event Processing systems is to efficiently detect complex patterns over streams of events. A CEP engine is responsible for filtering and combining primitive events into higher-level, *complex events*, which are then reported to end users. Areas in which CEP is widely applied include financial services [15], RFID-based inventory management [27], and electronic health record systems [21].

The patterns recognized by CEP systems are normally created using declarative specification languages. They are usually based on relational languages extended with additional operators, making it possible to define a wide range of pattern types, such as sequences, disjunctions and iterations. Basic operators can be combined into arbitrarily composite expressions. Filters, constraints, time windows, and mutual conditions between events can be applied. Examples of complex event specification languages include SASE [28], CQL [5], CEL [8] and CEDR [7].

Different patterns require different policies as to whether a primitive event selected for some pattern match can be considered again for future matches, a notion known as *event selection strategy*. In this paper we assume that, for any pattern, all possible matching combinations of events are requested to be detected. Consequently, a primitive event is allowed to participate in an unlimited number of matches. This selection strategy is called *skip-till-any-match*[3].

To illustrate the above notions, consider the following example.

Example 1 *A stock market monitoring application is requested to detect non-standard behavior of stock prices. For each stock identifier, a primitive event is generated when its price exceeds some predefined value. We would like to detect a complex event in which an irregularly high price of stock St1 is detected, followed by a high price of stock St2, which is also followed by a high price of stock St3, within a time window of one hour. High price is defined as exceeding a given threshold T.*

Now, consider an input stream a_1, a_2, b_1, b_2, c, where events a, b and c denote observations of high values of stocks *St1*, *St2* and *St3* respectively. Under the skip-till-any-match selection strategy, the matches to be returned are $\{a_1b_1c\}$, $\{a_2b_1c\}$, $\{a_1b_2c\}$ and $\{a_2b_2c\}$.

Various methods have been discussed for evaluating complex event patterns. One popular approach is to employ Non-deterministic Finite Automata (NFAs) [12, 15, 20, 24, 28]. A query is compiled into an NFA consisting of a set of states and conditional transitions between them. States are arranged in a pattern-specific topology. Transitions are triggered by the arrival of an appropriate event on the stream. Other mechanisms have also been proposed, such as trees [23], finite state machines [26], and detection graphs [4], to name a few.

An NFA for Example 1 is displayed in Figure 1. At each point in time, an instance of the state machine is maintained for every detected sub-match of the pattern. This instance is kept until a full match is detected. As an example,

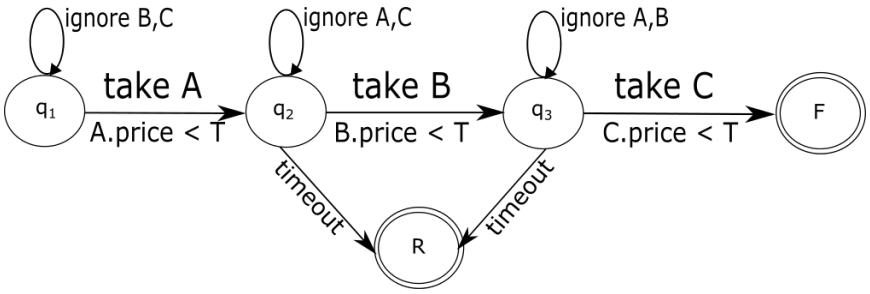

Figure 1: NFA for the pattern from Example 1

consider again a stream a_1, a_2, b_1, b_2, c. After the first three events have arrived, the system will maintain match prefixes $\{a_1\}$, $\{a_2\}$, $\{a_1 b_1\}$ and $\{a_2 b_1\}$. After the fourth event, two more instances will be added. Finally, following the arrival of c, the NFA detects four sequences matching the pattern.

This approach, however, can prove inefficient when events arrive at highly varied rates. Consider an input stream in which events of types A and B arrive at a rate of one per second, whereas an event of type C arrives once every 12 hours. In this case, the system must maintain a large number of prefixes that might not lead to any matches. Since the number of prefixes to be kept grows exponentially with the number of event types in a pattern, this method becomes highly wasteful in terms of memory and computational resources. The described situation could be avoided if the evaluation started from C, the rarest event type.

Several papers proposed to optimize the performance of CEP systems by modifying evaluation order and delaying the processing of more common events [4, 23, 26]. In [22], a *lazy evaluation* mechanism is described for sequence patterns. It constructs prefixes starting from the least frequent event, instead of the first event in the sequence. Partial matches are extended by adding events in ascending order of their frequency, rather than according to their order in the sequence. New types of NFAs, known as Chain NFAs and Tree NFAs, utilize the above principles. However, this method only applies to sequence patterns and lacks extensive experimental support.

In this paper, we generalize the lazy evaluation model to apply to most common types of patterns, including conjunctions, negations, iterations, and compositions. For some of them, we believe our work is the first to produce an efficient automata-based solution. The construction of a Lazy Chain NFA for each of the aforementioned pattern types is presented in detail. We also extend this topology for use with composite patterns. This extension, which we call a Lazy Multi-Chain NFA, is capable of detecting an arbitrary composition of the operators above. The correctness of all construction algorithms is formally proven. A comprehensive experimental evaluation over real-life data demonstrates a performance gain of two orders of magnitude over traditional NFA-based approaches, with significantly reduced memory resource requirements.

The remainder of the paper is organized as follows. Section 2 provides the required background and briefly describes the NFA evaluation framework. The concepts and ideas of lazy evaluation are outlined, accompanied by intuitive explanations and examples. We proceed to describe how a Lazy Chain NFA can be constructed using given frequencies of the participating events in Section 3. Lazy Multi-Chain NFA is discussed in Section 4. We formally prove the correctness of our construction in Section 5. Section 6 presents the experimental evaluation. Section 7 describes related work. Section 8 summarizes the paper.

2 Eager and Lazy Evaluation

In this section, we discuss in detail the two main parts of a CEP system: the specification language and the evaluation mechanism. For the former, SASE+[20] will be assumed for the rest of the paper. For the latter, we first present the "eager" NFA evaluation framework over SASE+ patterns as described in [3], processing every incoming event upon arrival. Then, the modifications required to employ the lazy evaluation principle are described. The SASE+ framework was chosen without loss of generality, and the same ideas may be applied to any NFA-based method.

2.1 Specification Language

The SASE+ language, thoroughly described in [20], combines a simple, SQL-like syntax with a high degree of expressiveness. The semantics and expressive power of the language are precisely defined in a formal model.

Each primitive event in SASE+ has an arrival timestamp, a type, and a set of attributes associated with the type. An attribute is a data item related to a given event type, represented by a name and a value.

In its most basic form, a complex event definition in SASE+ is composed of three building blocks: PATTERN, WHERE and WITHIN. The PATTERN clause defines the primitive events we would like to detect and the operator applied to combine them into a pattern match. Each event is represented by a unique name and a type. The WHERE clause specifies constraints on the values of data attributes of the primitive events participating in the pattern. These constraints may be combined using Boolean expressions. Finally, the WITHIN clause defines a time window over the entire pattern, specifying the maximal allowed time interval between the arrivals of the primitive events.

As an example, consider the pattern from Example 1. One possible representation of this complex event is:

$$
\begin{aligned}
&PATTERN\ SEQ\,(A\,a, B\,b, C\,c)\\
&WHERE\ skip_till_any_match\,\{\\
&\quad a.price > T\\
&\quad AND\ b.price > T\\
&\quad AND\ c.price > T\}\\
&WITHIN\ 1\ hour.
\end{aligned}
\tag{1}
$$

Here, we define three event types for stocks with identifiers A, B and C. Every primitive event represents a value of a respective stock at some point in time. Since a fixed order on stock price reports is defined, the sequence (SEQ) operator is to be used. We assume that a numerical attribute *price* is assigned to each event. *skip_till_any_match* denotes the event selection strategy as presented in Section 1.

2.2 The Eager Evaluation Mechanism

After a SASE+ pattern is created, it is compiled into an NFA, which is then employed on an input stream to detect pattern occurrences. In this section we define the structure of the eager NFA, which is a slightly modified version of the one described in [3].

Formally, an NFA is defined as follows:

$$A = (Q, E, q_1, F, R)\,,$$

where Q is a set of states; E is a set of directed edges; q_1 is an initial state; F is an accepting state; and R is a rejecting state.

Evaluation starts at the initial state. Transitions between states are triggered by event arrivals from the input stream. Each NFA instance is associated with a *match buffer*. As we proceed through an automaton towards the final state, we use the match buffer to store the events that caused the transitions. It is always empty at q_1, and events are gradually added to it during evaluation. This is done by executing actions of the traversed edges, as will be described shortly.

If during the traversal of an NFA instance the accepting state is reached, the content of the associated match buffer is returned as a successful match for the pattern. If the rejecting state is reached, the NFA instance and its match buffer are discarded. If the time window specified in the WITHIN block is violated, a special *timeout* event is generated, resulting also in the rejecting state.

An edge is defined by the following tuple:

$$e = (q_s, q_d, action, types, condition)\,,$$

where q_s is the source state of an edge; q_d is the destination state; *action* is always one of those described below; *type* may be one or more event types specified in the PATTERN block; and *condition* is a Boolean predicate, reflecting the conditions in the WHERE block that have to be satisfied for the transition to occur.

The *action* associated with an edge is performed when the edge is traversed. It can be one of the following:

- *take* - consumes the event from the input stream and adds it to the match buffer. A new instance is created.
- *ignore* - skips the event (consumes the event and discards it instead of storing). No new instance is created.

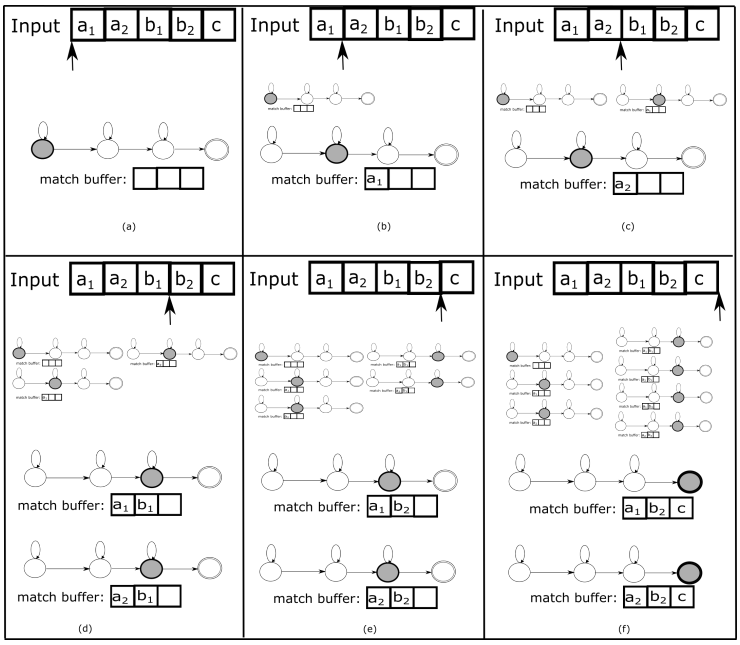

Figure 2: Example of a non-deterministic evaluation of the NFA from Figure 3. The rejecting state is omitted for simplicity. At each step, newly created instances are magnified. The current state of each instance is highlighted in gray.

As an example, consider again Figure 1. Note that the accepting state can only be reached by executing three *take* actions; hence, successful evaluation will produce a match buffer containing three primitive events which comprise the detected match.

Several *take* edges may lead from the same state and specify the same event type. In this case, an event will non-deterministically cause multiple traversals from a given state. If an event triggered n edges with a *take* action, the instance will be replicated n times, for each of the possible traversals.

Upon startup, the system creates a single instance with an empty match buffer, whose current state is the initial state. As events arrive, multiple instances of an NFA are created and run in parallel, one for each partial match detected up to that point. Every event received on the input stream will be applied to all NFA instances.

Figure 2 illustrates the eager evaluation process of the NFA from Figure 1, applied on an input stream a_1, a_2, b_1, b_2, c.

2.3 The Lazy Evaluation Mechanism

Here we briefly present the lazy evaluation mechanism, described fully in [22]. As demonstrated in Section 1, more frequent primitive events might trigger creation of instances that do not lead to any matches. Consider, for example, how the evaluation in Figure 2 might look if the input stream contained 100 events of type A, followed by 100 events of type B. Since no event of type C is present, no pattern match can exist. However, 100^2 instances would be spawned and redundant computations would take place. This could have been avoided simply by modifying the evaluation order to start from C.

There are two main reasons for the problem described above. First, an eager NFA is constructed in a manner that preserves the structure of the underlying pattern. Second, any primitive event must be processed immediately upon arrival. Hence, e.g., in a sequence pattern, re-ordering events with mutual temporal constraints is not possible.

The lazy evaluation model takes advantage of the varying arrival rates of the events in the sequence to significantly reduce the use of computational and memory resources. The core idea is to arrange NFA states according to the order of event frequencies, rather than in the way suggested by the pattern. We assume for now that all frequencies are known in advance. In Section 3 we will demonstrate how to execute this construction for the most common pattern types.

To achieve out-of-order evaluation, lazy NFA instances are designed to store incoming events and retrieve them later for processing. To that end, an additional buffer, referred to as the *input buffer*, is associated with each NFA instance. In addition, a new action, denoted as *store*, is defined. When an edge with a *store* action is traversed, the event causing the traversal is inserted into the input buffer. Expired events are silently removed from the buffer as the time window advances.

Retrieval from the input buffer is implemented by extending the semantics of the *take* action. Whereas in the eager NFA model an event accepted by this type of edge is always taken from the input stream, here it also triggers a search for events of the required type inside the input buffer. The results of this search, as well as events from the input stream, are then evaluated non-deterministically by spawning new NFA instances. If the search yields no result, or if the retrieved events fail to trigger a *take* transition due to unsatisfied conditions, a special event called *search_failed* is generated for the current instance.

Consider again the pattern from Example 1. Assume that the event type C is the least frequent, and the event type A is the most common. Figure 3 shows a possible lazy NFA, which detects the pattern in order *C,B,A*. Upon arrival of A or B, an outgoing *store* edge of q_1 inserts them into the input buffer. Following an arrival of type C, an outgoing *take* edge of q_2 retrieves events belonging to type B from the input buffer and attempts to add each of them to a partial match, as if they arrived from the input stream. The same occurs for events of type A in q_3.

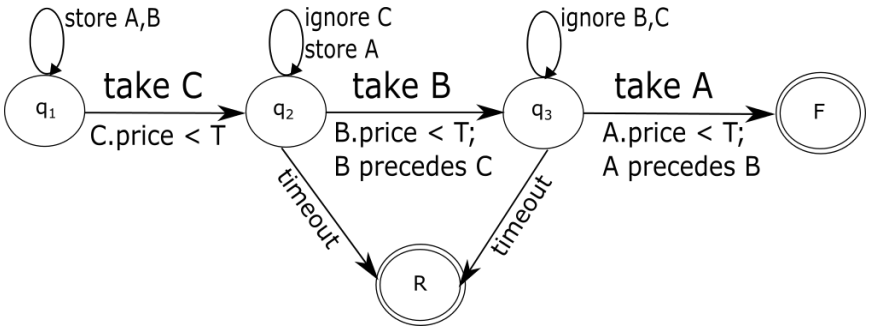

Figure 3: Lazy Chain NFA for Example 1 with frequency order *C,B,A*.

Note that invoking a full scan of the entire input buffer on each *take* action would be inefficient and unnecessary. While searching for events of type A, we are only interested in those that arrived before the already accepted B. In general, only a certain, usually very narrow range of events in the input buffer is relevant to a given *take* edge. Since the input buffer is arranged chronologically, these temporal constraints make it possible to efficiently limit the search to only a small portion of the storage.

To implement the desired functionality, we extend the NFA edge definition:

$$e = (q_s, q_d, action, types, condition, prec, succ)\,.$$

Here, *prec* and *succ*, also known as *ordering filters*, are sets of event types which enforce temporal limitations on an event taken by this edge. Elements in *prec* must precede this event, while elements in *succ* must succeed it. Both sets may only contain event types selected among those already taken during evaluation until this point. In the example above, for the edge taking A, the following holds: $prec\,(A) = \varnothing; succ\,(A) = \{B, C\}$. Thus, while attempting to traverse this edge, the system will scan the buffer from the beginning, but only until the timestamp of B is reached. For an event type E to be accepted from the input stream, we require the condition $succ\,(E) = \varnothing$.

The following section will provide ordering filter definitions for different types of lazy NFAs.

3 Lazy Chain NFA Construction

The Lazy Chain NFA is an NFA topology which utilizes the constructs of the lazy evaluation model presented in Section 2.3. This section contains our main contribution, which is the universal, pattern-independent definition of the Lazy Chain NFA and its applications for detecting a wide range of pattern types.

We will start by providing the intuition behind the construction of a generic Lazy Chain NFA. All pattern types discussed in this section will share this common structure with only minor adaptations. The following sub-sections will focus on each type, providing formal and detailed definitions.

Given a pattern containing n primitive event types, a corresponding Lazy Chain NFA consists of $n + 2$ states. The first $n + 1$ states are arranged in a chain according to ascending frequency order of the events, which we assume to be given in advance. Each of the first n states is responsible for detecting one primitive event in the pattern (the initial state detects the rarest event, the next state detects the second rarest event, etc.). The final state in a chain is the accepting state F. In addition, the rejecting state R is connected to all states except for q_1 and F. Its purpose is to collect invalid and expired partial matches.

We will denote by *freq* the ascending frequency order in which the first n states are arranged. We will also denote by e_i the i^{th} event type in *freq* and by q_i the corresponding state in the chain. The state q_i will generally have several types of outgoing edges. A *take* edge attempts to add the next event in the pattern to the partial match and to advance to the next state q_{i+1}. A *store* edge adds all events of types yet to be processed (succeeding e_i in *freq*) to the input buffer. An *ignore* edge discards events of already processed types (preceding e_i in *freq*). Finally, a *timeout* edge accepts the special *timeout* event and proceeds to R.

Figure 3 exemplifies the common structure of a Lazy Chain NFA. Note that it is based solely on relative frequencies of primitive events. It depends neither on the type of the input pattern nor on its structure (e.g., on the requested temporal sequence order). Instead, pattern-specific requirements will be expressed by the parameters of the edges, as will be explained in the subsequent sections.

3.1 Sequences

Sequences are patterns requiring a number of primitive events to arrive in a predefined order. Pattern 1 presented in Section 2.1 is an example of a sequence of three events. A Lazy Chain NFA for sequence patterns conforms fully to the common structure. Since temporal constraints are crucial for this pattern type, properly defining ordering filters on edges is of particular importance.

We will proceed now to the formal description. Let E_i denote the set of outgoing edges of q_i. Let $Prec_{ord}(e)$ denote all events preceding an event e in an order *ord*. Similarly, let $Succ_{ord}(e)$ denote all events succeeding e in *ord*. Then, E_i will contain the following edges:

- $e_i^{ignore} = (q_i, q_i, ignore, Prec_{sel}(e_i), true, \varnothing, \varnothing)$: any event whose type corresponds to one of the already taken events is ignored.

- $e_i^{store} = (q_i, q_i, store, Succ_{sel}(e_i), true, \varnothing, \varnothing)$: any event that might be taken in one of the following states is stored in the input buffer.
- $e_i^{timeout} = (q_i, R, ignore, timeout, true, \varnothing, \varnothing)$: if a timeout event is detected, the NFA instance proceeds to the rejecting state and is subsequently discarded.
- $e_i^{take} = (q_i, q_{i+1}, take, e_i, cond_i, prec_i, succ_i)$: an event of type e_i is taken only if it satisfies the conditions required by the initial pattern (denoted by $cond_i$).

Now we will define how ordering filters $prec_i$ and $succ_i$ are calculated. Given a set S of events, let $Latest(S)$ be the latest event in S (i.e., an event with the largest timestamp value). Correspondingly, let $Earliest(S)$ denote the earliest event in S. Finally, let seq denote the original sequence order as specified by the input pattern.

The ordering filters for a *take* edge e_i^{take} will be defined as follows:

$$prec_i = \begin{cases} Latest(Prec_{freq}(e_i) \cap Prec_{seq}(e_i)) & if\ Prec_{freq}(e_i) \cap Prec_{seq}(e_i) \neq \emptyset \\ \emptyset & otherwise \end{cases};$$

$$succ_i = \begin{cases} Earliest(Prec_{freq}(e_i) \cap Succ_{seq}(e_i)) & if\ Prec_{freq}(e_i) \cap Succ_{seq}(e_i) \neq \emptyset \\ \emptyset & otherwise \end{cases}.$$

It is easy to see that $prec_i$ and $succ_i$ consist of a single element each. $prec_i$ contains the latest event preceding e_i in the original sequence order, out of those already accepted when the state q_i is reached. Similarly, $succ_i$ holds the earliest event, out of those available, succeeding e_i in seq.

The Lazy Chain NFA for sequence patterns will thus be defined as follows:

$$A = (Q, E, q_1, F, R);$$
$$Q = \{q_i | 1 \leq i \leq n\} \cup \{F, R\};\ E = \bigcup_{i=1}^{n} E_i.$$

Figure 3 demonstrates the Lazy Chain NFA for the pattern from Section 2.1.

3.2 Conjunctions

Conjunction patterns detect a set of events in the input stream, regardless of their mutual order of arrival. This pattern type presents a considerable challenge to the traditional NFA-based approaches. The complication follows from the nature of a finite automaton, which requires specifying an order in which pattern elements will be accepted. However, as all orders between the participating events are valid, the only legitimate way to construct an NFA for the desired pattern is by incorporating all of them.

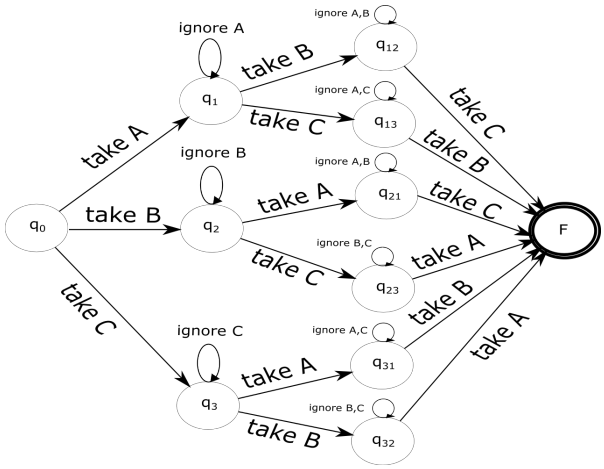

Figure 4: Eager NFA for Conjunction Pattern *AND(A,B,C)*. The rejecting state is omitted for simplicity.

Consider the following simple conjunction pattern:

$$\begin{aligned} &PATTERN\ AND\,(A\,a, B\,b, C\,c) \\ &WITHIN\,1\,hour. \end{aligned} \tag{2}$$

Figure 4 illustrates an eager NFA accepting this pattern. It can be seen that the number of states and transitions is exponential in length of the AND clause. As a result, the matching process becomes highly inefficient, even for a small number of event types.

The Lazy Chain NFA solves the aforementioned problem. Instead of attempting to match all possible orders, only the ascending frequency order is incorporated, as presented in Figure 3. To the best of our knowledge, this is the first method for automata-based detection of conjunction patterns that requires only a linear number of states.

The definition of the Lazy Chain NFA for conjunctions is similar to the one presented above for sequence patterns. The only difference is the absence of ordering filters. Since no constraints can be defined on the mutual order of primitive events, the entire content of the input buffer has to be examined during each search operation, with no possibility to discard parts of it. Thus, buffer searches are rather slow as compared to searches that take place during evaluation of sequences. However, this solution still significantly outperforms the traditional (eager) approach, as our experimental results in Section 6 demonstrate.

More formally, a *take* edge e_i^{take} for an event type e_i, detected by a state q_i, is defined as follows:

$$e_i^{take} = (q_i, q_{i+1}, take, e_i, cond_i, , \varnothing, \varnothing)\,.$$

Otherwise, the construction is identical to the one shown in Section 3.1.

3.3 Partial Sequences

Partial sequence patterns are conjunctions in which temporal constraints exist between subsets of the primitive events involved. As an example, consider a pattern:

$$\begin{array}{l} PATTERN\ AND\left(SEQ\left(A\ a, B\ b\right), SEQ\left(C\ c, D\ d\right), E\ e\right) \\ WITHIN\ 1\ hour. \end{array} \tag{3}$$

Here, an event of type B must appear after an event of type A. However, it may arrive either before or after events of types C, D and E. Similarly, D has to appear after C, and E can appear at any place in a match.

Note that sequence and conjunction patterns, as presented above, are two opposite edge cases of a partial sequence.

The Lazy Chain NFA for a partial sequence will be defined identically to the case of a full sequence. It will, however, incorporate ordering filters only for those events participating in at least one sub-sequence. Also, for events appearing in multiple sub-sequences, $prec_i$ and $succ_i$ will contain several values. Among those, only the most restrictive ones will be chosen at runtime.

We will now proceed to the formal definition of ordering filters. Let $SEQ = \{seq_1, \cdots, seq_k\}$ denote all sub-sequences in a pattern. Note that the sub-sequences in SEQ do not necessarily contain independent sets. In addition, let

$$Prec_{SEQ}\left(e\right) = \bigcup_{seq \in SEQ} Prec_{seq}\left(e\right)$$

$$Succ_{SEQ}\left(e\right) = \bigcup_{seq \in SEQ} Succ_{seq}\left(e\right).$$

The ordering filters for an edge e_i^{take} are then defined as follows:

$$prec_i = \begin{cases} Latest\left(Prec_{freq}\left(e_i\right) \cap Prec_{SEQ}\left(e_i\right)\right) & if\ Prec_{freq}\left(e_i\right) \cap Prec_{SEQ}\left(e_i\right) \neq \emptyset \\ \emptyset & otherwise \end{cases};$$

$$succ_i = \begin{cases} Earliest\left(Prec_{freq}\left(e_i\right) \cap Succ_{SEQ}\left(e_i\right)\right) & if\ Prec_{freq}\left(e_i\right) \cap Succ_{SEQ}\left(e_i\right) \neq \emptyset \\ \emptyset & otherwise \end{cases}.$$

Otherwise, the construction is identical to the one shown in Section 3.1.

3.4 Negations

In a pattern with negation, some of the primitive event types are not allowed to appear at the predefined places. We will denote them as *negated events*. Negated events can be specified anywhere in a pattern and form mutual conditions with positive events. Patterns of any type may include a negation part.

The following is an example of a sequence pattern with a negated event:

$$\begin{array}{l} PATTERN\ SEQ\left(A\ a, NOT\left(B\ b\right), C\ c, D\ d\right) \\ WHERE\ skip_till_any_match\ \{b.x < c.y\} \\ WITHIN\ 1\ hour. \end{array} \tag{4}$$

In this case, a successful pattern match will contain instances of A, C and D alone. Note that only events of type B satisfying the condition with an event of type C are not allowed to appear.

Existing NFA-based CEP frameworks employ two different techniques for treating negated events. The first is to check for negative conditions as a post-processing step, after the accepting state is reached. This technique introduces a potential performance issue. Consider a case in which events of type B are very frequent and their presence results in the discarding of all partial matches. Matches detected by NFA will thus be invalidated only during post-processing, causing superfluous computations. This situation could be avoided by moving the negated event check to an earlier stage.

The second technique consists of augmenting an NFA with "negative edges" that lead to a rejecting state upon detection of a forbidden event. To the best of our knowledge, existing solutions of this kind only solve limited cases. Namely, only sequences are considered and no conditions between primitive events are supported. The reason for these restrictions is that, when no event buffering is used, it is impossible to verify the absence of a negated event that depends on some future event. Consider again the example above. When an event of type B arrives, an eager NFA cannot check whether it satisfies a condition with some event of type C, since C has not yet been received. Our framework avoids this issue, as evaluation is performed out of order.

We propose two Lazy Chain NFA-based solutions for detecting negations. The first, which we call a Post-Processing Chain NFA, implements the post-processing paradigm outlined above. The second, First-Chance Chain NFA, attempts to detect a negated event as soon as possible. From an analytical standpoint, neither solution is superior to the other, and each can be favorable in different situations. A more efficient mechanism for a particular negation pattern can be easily selected, either automatically or manually.

3.4.1 Post-Processing Negation

The Lazy Post-Processing Chain NFA first detects a sub-chain of positive events, and then proceeds to a second sub-chain, where each state corresponds to a single negated event. For this negative sub-chain, descending frequency order is used (as opposed to ascending order for a positive part). Each *negative state* is responsible for verifying absence of one negated event. Thus, transitions between negative states are triggered by either reaching a timeout or by an unsuccessful search in the input buffer. These situations are indicated by special *timeout* and *search_failed* events. Arrival of a forbidden event in a negative state triggers a transition to the rejecting state. The last negative state is followed by the accepting state.

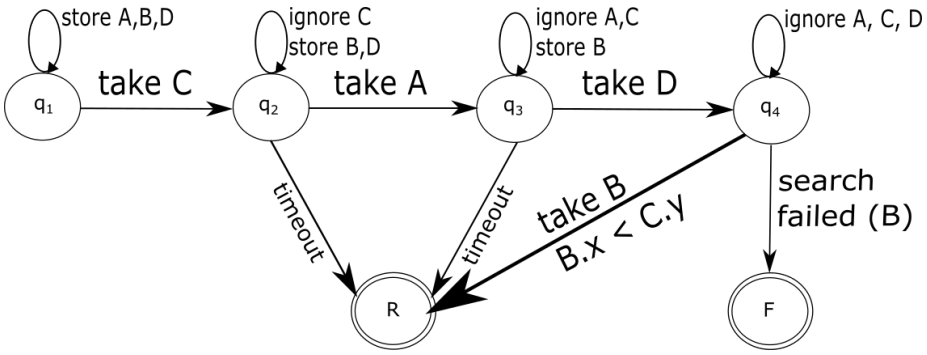

Figure 5: Lazy Post-Processing Chain NFA for Pattern 4. The frequency order of C,A,D is assumed. Ordering filters are omitted for clarity.

Figure 5 demonstrates the Post-Processing Chain NFA for Pattern 4. Since we only need to check for an occurrence of B before C, a scan of the input buffer is sufficient. Hence, only a *search_failed* event is expected. For patterns in which a negated event may appear at the end (e.g., for conjunctions), an edge taking *timeout* is also required.

We will now formally define the Lazy Post-Processing Chain NFA.

Let $P = \{e_1, \cdots, e_k\}$ be all positive event types in a pattern. Let $N = \{h_1, \cdots, h_l\}$ be all negated event types. We will denote by $freq_p$ the ascending frequency order of the events in P, and by $freq_n$ the descending frequency order of the events in N. Let $Q_p = \{q_1, \cdots, q_k\}$ be a set of states corresponding to positive events ordered according to $freq_p$, and let $Q_n = \{r_1, \cdots, r_l\}$ be a set of states corresponding to positive events ordered according to $freq_n$. Finally, let E_q denote the set of outgoing edges of a state q.

For each positive state $q_i \in Q_p; i \leq k$, the edges in E_{q_i} are defined as for the Chain NFA for the underlying pattern, with the exception of an edge e_i^{store}, which stores all events in N in addition to those in $Succ_{freq_p}(e_i)$. Also, for q_k, the *take* edge e_i^{take} leads to the first negative state, r_1.

For a negative state $r_i \in Q_n; i \leq l$, the edges in E_{r_i} are:

- $e_i^{ignore} = (r_i, r_i, ignore, Prec_{freq_n}(e_i) \cup P, true, \varnothing, \varnothing)$: a positive event or previously checked negative event is ignored.
- $e_i^{store} = (r_i, r_i, store, Succ_{freq_n}(e_i), true, \varnothing, \varnothing)$: an event which may be potentially taken in one of the following states is stored in the input buffer.
- $e_i^{take} = (r_i, R, take, h_i, cond_i, prec_i, succ_i)$: an instance of a negated event h_i satisfying the conditions triggers a transition to the rejecting state.

If h_i can only be accepted from the input buffer (i.e., the condition $succ(h_i) \neq \varnothing$ holds), an additional edge is present:

- $e_i^{search_failed} = (r_i, r_{i+1}, ignore, search_failed, true, \varnothing, \varnothing)$: in case of a failed search for a buffered event, the execution successfully proceeds to the next state.

Otherwise, if h_i can be accepted from the input stream, an additional edge is present:

- $e_i^{timeout} = (r_i, r_{i+1}, ignore, timeout, true, \varnothing, \varnothing)$: in case of a timeout, the negation test is passed and the execution proceeds to the next state.

Also, for r_l, the *timeout* or *search_failed* edge leads to the accepting state F. The ordering filters $prec_i$ and $succ_i$ are calculated in the same manner as for the underlying pattern type.

We are now ready to define the Lazy Post-Processing Chain NFA:

$$A = (Q, E, q_1, F, R)\,;$$
$$Q = Q_p \cup Q_n \cup \{F, R\}\,;\; E = \bigcup_{q \in Q_p \cup Q_n} E_q.$$

Although this NFA shares the previously discussed drawbacks of the post-processing method, it also has several benefits over the other negation NFA, described below. First, since event buffering is an inherent part of the Lazy Evaluation mechanism, implementing this NFA on top of the existing framework is straightforward. Second, in some scenarios the best strategy is to postpone negation until the end. One example is a very frequent negated event with a highly selective filter condition. Third, this is the only possible approach for conjunctions with negation and for sequences with a negated event at the end. In these cases, we have to wait until the time window expires to perform a negation check.

3.4.2 First-Chance Negation

The Lazy First-Chance Chain NFA implements a paradigm opposite to that of the Post-Processing Negation NFA. It operates by pushing the detection of negated events to the earliest point possible. The key observation is that it is often unnecessary to wait for all positive events to arrive before launching a negated event check.

Consider Pattern 4 again. Clearly, a potential event of type B is only dependent on events of types A (temporal condition) and C (explicit and temporal conditions). Consequently, we only need to have A and C in our match buffer to check for conflicting B events. However, a solution based on post-processing will also wait for D to arrive. As a result, if an event of type B is found, it will lead to discarded matches and to redundant operations, including the same search for B for every instance of D. Furthermore, if D had mutual conditions with A or C, they would be verified repeatedly, only to be invalidated later.

Lazy First-Chance Chain NFA overcomes this performance issue. For each negated event, it determines the earliest state in which a check for this event can be executed. From this state, a *take* edge is added, which leads to the rejecting

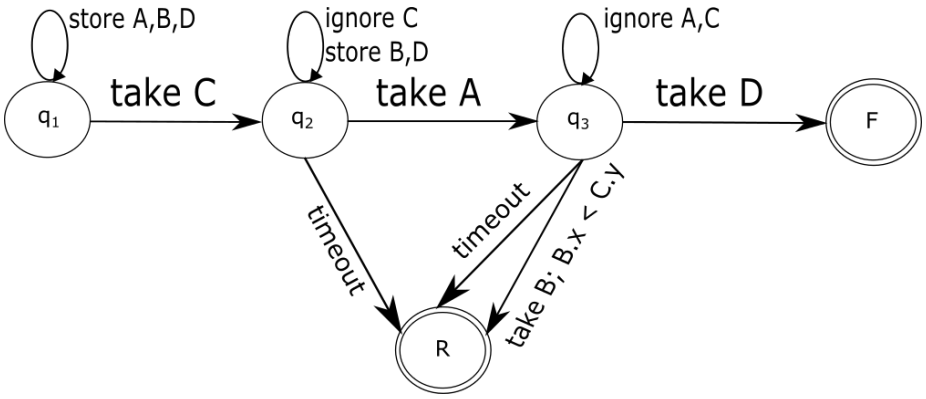

Figure 6: Lazy First-Chance Chain NFA for Pattern 4. The frequency order of C,A,D is assumed. For simplicity, ordering filters are omitted.

state if an event matching the conditions is encountered. This type of NFA therefore does not have a state corresponding to each negated event. Instead, it consists only of the positive chain, augmented with the aforementioned *take* edges.

As mentioned above, this approach cannot be efficiently applied on a pattern in which a negated event may appear after all positive events in a match. This is because for such an event, the earliest state in which it can be detected is the accepting state.

Figure 6 demonstrates the First-Chance Chain NFA for Pattern 4. As the figure shows, the absence of B is verified before D is accepted, which is indeed the earliest point possible for this pattern. Note that there are two edges between q_3 and F, one handling the timeout case and the other detecting the negated event type B.

We will now proceed to formally define the First-Chance Chain NFA.

Let $P = \{e_1, \cdots, e_k\}$ be all positive event types in a pattern and let $N = \{h_1, \cdots, h_l\}$ be all negated event types. Let $A_{pos} = (Q_{pos}, E_{pos}, q_1, F, R)$ denote a Lazy Chain NFA for the positive part of the pattern. For each negated event h_i we will define the following:

- $ImmPrec(h_i) = Latest(Prec_{SEQ}(h_i))$, the latest detected event preceding h_i.
- $ImmSucc(h_i) = Earliest(Succ_{SEQ}(h_i))$, the earliest detected event succeeding h_i.
- $Cond(h_i)$, set of all event types forming mutual conditions with h_i.
- $DEP(h_i) = \{ImmPrec(h_i), ImmSucc(h_i)\} \cup Cond(h_i)$, set of all event types which must be detected before the absence of h_i can be validated.

- $q_{DEP}(h_i)$, the earliest state in Q_{pos} in which all events in $DEP(h_i)$ are already detected.

The First-Chance Chain NFA will then be constructed by augmenting E_{pos} with an edge from $q_{DEP}(h_i)$ to R for each h_i:

$$A = (Q_{pos}, E_{pos} \cup E_{rej}, q_1, F, R);$$
$$E_{rej} = \{(q_{DEP}(h_i), R, take, h_i, cond_i) \,|h_i \in N\}.$$

In addition, each *store* edge in E_{pos} will be modified to apply to events in N as well.

3.5 Iterations

The term iteration operator (also called Kleene closure) refers to patterns in which given events are allowed to appear multiple and unbounded numbers of times. Detecting iterations is particularly challenging under the skip-till-any-match event selection strategy because of the exponential number of output combinations [31]. Our solution utilizes lazy evaluation principles to minimize the number of NFA instances and the calculations performed during the detection process.

For clarity of presentation, we will only discuss sequence patterns with a single iterated event. The concepts presented below can be easily extended to conjunction and partial sequence patterns and to an arbitrary number of events under iteration.

Consider the following sequence pattern with an iterated event:

$$\begin{array}{l} PATTERN\ SEQ(A\ a, B+\ b[], C\ c) \\ WITHIN\ 1\ hour. \end{array} \tag{5}$$

For a sample input stream a, b_1, b_2, b_3, c, the expected output will be:

$$ab_1c, ab_2c, ab_3c, ab_1b_2c, ab_1b_3c, ab_2b_3c, ab_1b_2b_3c.$$

Our approach is to convert this pattern into a regular sequence $SEQ(A, B, C)$ and to address the subsets of b_1, b_2, b_3 as separate "events". For example, the input stream above can be converted to:

$$a(b_1)(b_2)(b_3)(b_1b_2)(b_1b_3)(b_2b_3)(b_1b_2b_3)c.$$

In this case our artificial "$B+$" event type will become the most frequent regardless of the original frequency of B. Consequently, if we were to construct an ordinary Chain NFA for this sequence, we would place the state responsible for detecting $B+$ at the end, as this event type would be the last in the ascending frequency order.

Following this principle, two modifications are required for constructing an Iteration Chain NFA from a Sequence Chain NFA. First, an iterated event type has to be placed at the end of the frequency order, regardless of its actual arrival

rate. Second, the *take* edge corresponding to this event is required to produce all subsets of its instances.

To implement the second modification, we introduce a new edge action called *iterate*. An *iterate* edge operates similarly to a *take* edge, consuming an event and adding it to the match buffer. However, as an *iterate* edge traverses the input buffer, it produces subsets of events belonging to the required type and returns them all. For example, if an NFA with an *iterate* edge for type B is applied on a stream above, an input buffer will contain 3 B events and the *iterate* edge will return 7 subsets, thus detecting 7 matches for the pattern.

When an event can be taken from the input stream, an *iterate* edge will add the new event to the input buffer, and then will only generate subsets including this event.

A Lazy Chain NFA for Pattern 5 looks identical to the one displayed in Figure 3 for Pattern 1, with an event type B pushed to the end and its corresponding edge action changed to *iterate*.

Now we are ready to formally define a Lazy Iteration Chain NFA. Let our pattern be $P = SEQ\,(e_1, \cdots, e_k*, \cdots, e_n)$ and let $freq = e_{i_1}, \cdots, e_k, \cdots, e_{i_n}$ denote the ascending frequency order of the primitive event types. The desired automaton will be created by the following steps:

1. Define $freq' = e_{i_1}, \cdots, e_{i_n}, e_k$ (an order identical to *freq* except for moving e_k to the end).
2. Create a Lazy Chain NFA A_{seq} for $P' = SEQ\,(e_1, \cdots, e_k, \cdots, e_n)$ with respect to the order $freq'$.
3. Let $e_k^{take} = (q_{e_k}, F, take, e_k, cond_{e_k}, prec_k, succ_k)$ denote the *take* edge for e_k in A_{seq}. Define $e_k^{iterate} = (q_{e_k}, F, iterate, e_k, cond_{e_k}, prec_k, succ_k)$.
4. Produce a new NFA $A_{iterate}$ by replacing e_k^{take} in A_{seq} with $e_k^{iterate}$.

More formally, if $A_{seq} = (Q, E, q_1, F, R)$, then:

$$A_{iterate} = \left(Q, (E \setminus \left\{e_k^{take}\right\}) \cup \left\{e_k^{iterate}\right\}, q_1, F, R\right).$$

3.5.1 Aggregations

Aggregation functions (SUM. AVG, MIN, MAX, etc.) can easily be integrated into the framework using the presented approach for iteration patterns. Since the lazy evaluation mechanism inherently supports event buffering, aggregation is performed in a straightforward manner by invoking a desired function on the input buffer contents. Note that an aggregation function can only be applied on events under an iteration operator. Consider the following example:

$$\begin{array}{l} PATTERN\ SEQ\,(A\,a, B+\ b[], C\,c) \\ WHERE\ skip_till_any_match\,\{AVG\,(b\,[i]\,.x) < c.y\} \\ WITHIN\ 1\ hour. \end{array} \tag{6}$$

The condition will be evaluated upon traversal of a corresponding *iterate* edge. For each fetched subset of the available B events, the edge action will apply an aggregate function on this subset and validate the condition.

3.5.2 Repetitions

A repetition operator is a bounded version of an iteration. In repetition patterns, we require a primitive event to appear at least l and at most m times. An iteration pattern is thus a repetition pattern with l set to one and m set to infinity.

Support for repetitions can be trivially added by defining lower and upper bounds on the size of a subset returned by an *iterate* edge.

3.5.3 Group-By-Attribute Optimization

In some cases, creating all subsets of events before any processing takes place is not the most efficient strategy. For example, consider the following iteration pattern:

$$\begin{aligned} &PATTERN\ SEQ\,(A\,a, B+\ b[], C\,c) \\ &WHERE\ skip_till_any_match\,\{b\,[i]\,.x = b\,[i-1]\,.x\} \\ &WITHIN\ 1\ hour. \end{aligned} \tag{7}$$

For this pattern, most subsets of instances of B will not be valid, since all subset members are required to share the same value of the attribute x. Therefore, the evaluation procedure described above will perform poorly, even compared to an eager approach.

For scenarios like the one shown, our framework introduces an optimization that allows a user to specify a *group-by-attribute*. Instances of the iterated event are then hashed in the input buffer by the value of this attribute. Upon traversal of an iterate edge, the generated subsets will only contain events sharing the same attribute value.

3.6 Lazy Tree NFA

A major drawback of the Lazy Chain NFA is its explicit requirement to know the frequencies of primitive events in advance. In many real-world scenarios, this information may not be available or may be subject to constant changes. In [22], the authors present Lazy Tree NFA, an alternative NFA topology for sequence pattern detection. A Lazy Tree NFA also employs lazy evaluation; however, it does not have the aforementioned precondition. Instead, it computes the evaluation order at each step in an ad hoc manner.

Since Lazy Tree NFAs and Lazy Chain NFAs are constructed using the same principles, the construction method can be easily applied to all the pattern types examined above.

4 Lazy Multi-Chain NFA

Despite the flexibility of the Lazy Chain NFA, its simple structure is insufficient for detecting more complex patterns, involving separate sets of events and matching conditions. One notable example is a disjunction. Disjunction patterns consist of several sub-patterns, of which only one needs to be detected for the successful match to be reported. Consequently, an NFA detecting a disjunction must include multiple paths from the initial to the final state. However, a Lazy Chain NFA only contains a single path between q_1 and F, which renders it unsuitable for processing patterns of this type.

We will overcome this limitation by defining a new topology, which we call *Lazy Multi-Chain NFA*. This type of NFA will possess all the qualities of the Lazy Chain NFA, but will also provide path selection functionality. To construct it, we will first produce a Lazy Chain NFA for each of the nested pattern parts. Then, we will merge the initial, the accepting, and the rejecting states of all sub-automata and join all the remaining states into a single automaton. The resulting NFA will thus be structured as a union of chains leading to the accepting state. Whenever a match for a sub-query is retrieved, one of the paths will be traversed and the match will be reported. Figure 7 demonstrates an example of this construction.

In Section 4.1 we will formally describe the structure of a Lazy Multi-Chain NFA for disjunction patterns. We will then present a generalization to composite patterns in Section 4.2.

4.1 Disjunctions

Disjunctions are the most commonly used type of composite patterns. They consist of multiple sub-patterns, which can belong to either of the types discussed above. As a set of events satisfying at least a single sub-pattern is detected, it is reported as a match for the whole pattern. For example, consider the following:

$$\begin{array}{l} PATTERN\ OR\left(SEQ\left(A\ a, B\ b, C\ c\right), SEQ\left(C\ c, D\ d, E\ e\right)\right) \\ WITHIN\ 1\ hour. \end{array} \tag{8}$$

A Lazy Multi-Chain NFA for this pattern is displayed in Figure 7. Note that a single primitive event is allowed to appear in multiple sub-patterns. In this case, upon arrival of an event of type C, two edge traversals from q_1 will be executed in a non-deterministic manner and two new NFA instances will be created.

We will now proceed to the formal definition. Let $p_1, \cdots, p_m$ be the sub-patterns of the disjunction pattern and let

$$A_1 = \left(Q_1, E_1, q_1^1, F_1, R_1\right)$$

$$\vdots$$

$$A_m = \left(Q_m, E_m, q_1^m, F_m, R_m\right)$$

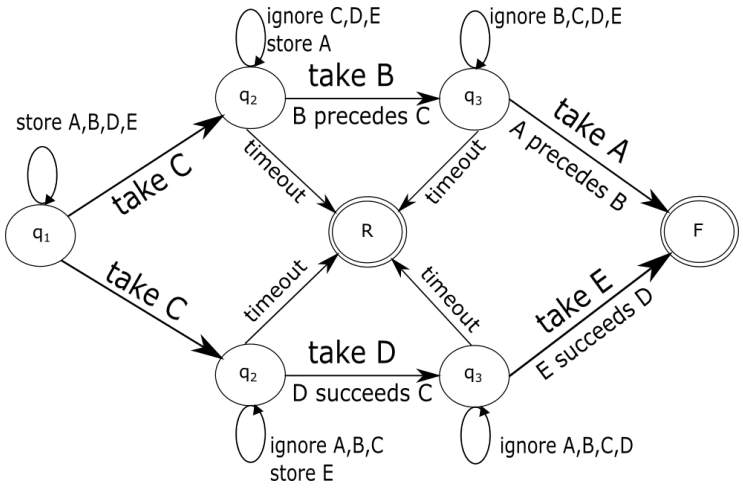

Figure 7: Lazy Multi-Chain NFA for Pattern 8 with frequency order *C,B,A,D,E*.

denote the Lazy Chain NFA for $p_1, \cdots, p_m$.

Let q_1, F, R be the initial, the accepting, and the rejecting states of the new Multi-Chain NFA A_{OR}, respectively.

The following definitions describe the outgoing edges of the initial states of $A_1, \cdots, A_m$ and the incoming edges of their final (accepting and rejecting) states.

$$E^j_{start} = \left\{e | e = \left(q^j_1, r, action, type, condition\right)\right\};$$

$$E^j_{acc} = \{e | e = (q, F_j, action, type, condition)\};$$

$$E^j_{rej} = \{e | e = (q, R_j, action, type, condition)\};$$

$$E_{start} = \bigcup_{j=1}^{m} E^j_{start};\ E_{acc} = \bigcup_{j=1}^{m} E^j_{acc};\ E_{rej} = \bigcup_{j=1}^{m} E^j_{rej}.$$

Now, we will define the new edges that will replace the existing ones and lead to q_1, F, R in the new NFA.

$$E^j_{OR-start} = \{e | e = (q_1, r, action, type, condition)\};$$

$$E^j_{OR-acc} = \{e | e = (q, F, action, type, condition)\};$$

$$E^j_{OR-rej} = \{e | e = (q, R, action, type, condition)\};$$

$$E_{OR-start} = \bigcup_{j=1}^{m} E^j_{OR-start};$$

$$E_{OR-acc} = \bigcup_{j=1}^{m} E^j_{OR-acc};$$

$$E_{OR-rej} = \bigcup_{j=1}^{m} E^j_{OR-rej}.$$

In addition, let

$$Q_{start} = \left\{ q_1^j | 1 \leq j \leq m \right\};$$
$$Q_F = \{ F_j | 1 \leq j \leq m \};$$
$$Q_R = \{ R_j | 1 \leq j \leq m \}.$$

Now, we are ready to define the Multi-Chain NFA:

$$A_{OR} = (Q, E, q_1, F, R);$$
$$Q = \left(\left(\bigcup_{j=1}^{m} Q_j \right) \setminus (Q_{start} \cup Q_F \cup Q_R) \right) \cup \{ q_1, F, R \};$$
$$E = \left(\left(\bigcup_{j=1}^{m} E_j \right) \setminus (E_{start} \cup E_{acc} \cup E_{rej}) \right) \cup (E_{OR-start} \cup E_{OR-acc} \cup E_{OR-rej}).$$

4.2 General Composite Patterns

A disjunction pattern can be alternatively viewed as a Boolean formula normalized to its DNF form. Moreover, all non-unary operators presented earlier can be expressed as operations of Boolean calculus[1]. Since any Boolean statement can be converted to DNF, it is possible to employ a Lazy Multi-Chain NFA for matching an arbitrary composite SASE+ pattern. The only step to be added to the construction algorithm from the previous section is transforming the input pattern accordingly.

It should be noted, however, that applying the above procedure may cause some sub-expressions to appear in multiple branches of the main disjunction pattern. This will, consequently, lead to superfluous calculations and NFA instances. For example, consider the following pattern:

$$\begin{array}{l} PATTERN\ SEQ\,(A\,a, OR\,(B\,b, C\,c)\,, D\,d) \\ WHERE\ skip_till_any_match\,\{a.price > d.price\} \\ WITHIN\ 1\ hour. \end{array} \quad (9)$$

The DNF form of this pattern is

$$OR\,(SEQ\,(A\,a, B\,b, D\,d)\,, SEQ\,(A\,a, C\,c, D\,d))\,.$$

The derived Multi-Chain NFA will contain two branches, both of which may perform computations for same pairs of events of types A and D. The problem becomes even more severe if B and C are negated types. In this case, the same output match <*a,d*> may be reported twice, since it satisfies both sub-patterns.

This issue may be solved by applying known multi-query techniques, e.g., as described in [15], and is a subject of our future work.

[1]For the purpose of the DNF conversion procedure, we represent a sequence pattern (full or partial) as a conjunction with additional temporal conditions between primitive events.

5 Equivalence of the Eager and the Lazy Chain NFA

In this section, we formally prove the correctness of the Lazy Chain NFA construction. It is shown that, for any pattern over the presented operators, a Lazy Chain NFA or a Lazy Multi-Chain NFA constructed according to Sections 3 and 4 is equivalent to the corresponding eager NFA in terms of the language it accepts.

Theorem 1 *Let P be a complex event pattern, as defined in Section 2.1, over the event types $e_1, \cdots, e_n$. Let $A_{eager} = (Q_{eager}, E_{eager}, q_1^{eager}, F_{eager}, R_{eager})$ be an eager NFA derived from P as described in Section 2.2 and let $A_{lazy} = (Q_{lazy}, E_{lazy}, q_1^{lazy}, F_{lazy}, R_{lazy})$ be a Lazy Chain NFA or a Lazy Multi-Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent, i.e., A_{eager} accepts a stream of events if and only if A_{lazy} accepts it.*

We will start with the outline of the proof. It will proceed in several steps.

First, we examine the case where P is a pure sequence pattern. We show that swapping two adjacent states in the NFA detecting P does not affect its correctness when the input buffer is used. From there, the statement is proven by induction for an arbitrary detection order. Next, the case of a sequence with negation is shown separately for both lazy mechanisms introduced above. In both parts, NFA construction properties are used to infer correctness. We prove the case of a sequence with iteration by induction on the length of an iterated pattern match, viewing the pattern as a regular sequence during the induction step.

Afterwards, we continue to more complex pattern types. We start by observing that conjunctions and partial sequences can be represented as disjunctions of full sequences. Then, we prove that two branch-structured NFAs (i.e., built of several chains starting at the initial state) are equivalent if their sub-automata are equivalent. Similarly, we demonstrate that a branch-structured NFA is equivalent to a chain-structured NFA, if all of the branches are equivalent to this chain-structured NFA. On the basis of these three statements, we derive the correctness for conjunctions and partial sequences as well. From there, the case of disjunctions is shown in a similar manner. We complete the proof by generalizing this case for an arbitrary composite pattern, using DNF form conversion.

Lemma 1 *Let P be a sequence pattern over $e_1, \cdots, e_n$, let seq be the order of the events in P, and let $freq$ be an ascending frequency order. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

We first address the case where the order $freq$ is identical to seq, i.e., we show that a lazy NFA for P where the frequency order is identical to the sequence order is equivalent to the eager NFA for P. Next we show that if $freq'$ is an

order received by swapping the event types in location i and $i+1$ ($1 \leq i \leq n-1$) in $freq$, then A_{lazy} is equivalent to $A_{lazy'}$, which uses $freq'$ instead of $freq$. This will conclude the proof, since we can obtain $freq$ by a set of swaps of adjacent event types on seq.

For the case $seq = freq$, the ordering filters of any edge will point to the end of the input buffer, and all events will be taken from the input stream; hence, the conditions on edges will become the same as in the eager NFA. In addition, the order of states will be the same as in the eager NFA. Consequently, the transition function between states is the same, making the two automata identical.

Now assume that $freq'$ is received by swapping the event types at location i and $i+1$. Let $A_{lazy'}$ be a corresponding Lazy Chain NFA derived from P and $freq'$.

Note that A_{lazy} and $A_{lazy'}$ differ only in the definitions of the outbound edges (including the self-loops) from the i^{th} and $(i+1)^{th}$ states, in particular in the ordering filters on their *take* edges. Let $s_1, ..., s_m$ be a stream of events. We examine the processing of this stream by both A_{lazy} and $A_{lazy'}$. We focus on instances of A_{lazy} that have reached the state q_{i+2}^{lazy}. We show that for every such instance, there is a corresponding instance in $A_{lazy'}$ that has reached $q_{i+2}^{lazy'}$ with the same input and match buffers. It is sufficient to show one direction (that instances in A_{lazy} have corresponding instances in $A_{lazy'}$), since the proof for the other direction is symmetric. Proving such a mapping will show that A_{lazy} and $A_{lazy'}$ are equivalent, since the processing of these instances will be identical in both NFAs from this point on.

Let us examine an instance I that has reached q_{i+2}^{lazy} in A_{lazy}. We will denote the event that caused I to transition into q_{i+2}^{lazy} by s_l. In addition, let us denote the event that caused I to transition into q_i^{lazy} by s_k. Let t_k denote the arrival time of s_k. Since the definitions of A_{lazy} and $A_{lazy'}$ are identical up to the i^{th} state, at time t_k there is an instance I' of $A_{lazy'}$ with match and input buffers identical to those of I. Furthermore, s_k will also cause I' to transition into $q_i^{lazy'}$. From this point on, the processing of the two instances may diverge, but we show that after the event s_l is processed, both instances will have transitioned into q_{i+2} with the same input and match buffers.

Any event corresponding to an event type already located in the match buffer will be ignored by both I and I'. Similarly, any event corresponding to an event type requested by one of the later states (q_{i+2} and beyond) will be stored to respective input buffers of both instances. Thus, we are left to consider the events that trigger the transitions $q_i^{lazy} \rightarrow q_{i+1}^{lazy}$ and $q_{i+1}^{lazy} \rightarrow q_{i+2}^{lazy}$.

We will examine the four possible scenarios for the above transitions in A_{lazy}:

1. The transition $q_i^{lazy} \rightarrow q_{i+1}^{lazy}$ occurred due to a corresponding event s_u arriving from the input stream, and the transition $q_{i+1}^{lazy} \rightarrow q_{i+2}^{lazy}$ occurred due to a corresponding event s_v taken from the input buffer. Then we may conclude that s_v arrived before s_u. Hence, during evaluation in $A_{lazy'}$, after time t_k the event s_v is either already located in the input buffer, or

will arrive eventually before s_u, triggering the transition $q_i^{lazy'} \rightarrow q_{i+1}^{lazy'}$, adding itself to the match buffer and removing itself from the input buffer if taken from there. Then, the event s_v will eventually be received from the input stream and trigger the transition $q_{i+1}^{lazy'} \rightarrow q_{i+2}^{lazy'}$, adding itself to the match buffer. Thus, at time t_l, both A_{lazy} and $A_{lazy'}$ contain s_u and s_v in their respective match buffers and do not contain s_v in their input buffers.

2. The transition $q_i^{lazy} \rightarrow q_{i+1}^{lazy}$ occurred due to a corresponding event s_u taken from the input buffer, and the transition $q_{i+1}^{lazy} \rightarrow q_{i+2}^{lazy}$ occurred due to a corresponding event s_v arriving from the input stream. Then, during evaluation in $A_{sel'}$, the event s_u is already located in the input buffer by the time s_v arrives from the input stream. When s_v arrives, it triggers the transition $q_i^{lazy'} \rightarrow q_{i+1}^{lazy'}$ and adds itself to the match buffer. Then, the event s_u will be immediately received from the input stream and will trigger the transition $q_{i+1}^{lazy'} \rightarrow q_{i+2}^{lazy'}$, adding itself to the match buffer and removing itself from the input buffer. Thus, at time t_l, both A_{lazy} and $A_{lazy'}$ contain s_u and s_v in their respective match buffers and do not contain s_u in their input buffers.

3. Both $q_i^{lazy} \rightarrow q_{i+1}^{lazy}$ and $q_{i+1}^{lazy} \rightarrow q_{i+2}^{lazy}$ were triggered by the corresponding events s_u, s_v arriving on the input stream. Then we may conclude that s_u arrived before s_v. During evaluation in $A_{lazy'}$, the event s_u will be inserted to the input buffer upon its arrival, which will occur after t_k. Later, the event s_v will trigger the transition $q_i^{lazy'} \rightarrow q_{i+1}^{lazy'}$, adding itself to the match buffer. Immediately afterwards, the event s_u located in the input buffer will trigger $q_{i+1}^{lazy'} \rightarrow q_{i+2}^{lazy'}$, adding itself to the match buffer and removing itself from the input buffer. Thus, at time t_l, both A_{lazy} and $A_{lazy'}$ contain s_u and s_v in their respective match buffers and do not contain s_u in their input buffers.

4. Both $q_i^{lazy} \rightarrow q_{i+1}^{lazy}$ and $q_{i+1}^{lazy} \rightarrow q_{i+2}^{lazy}$ were triggered by the corresponding events taken from the input buffer. Then, during evaluation in $A_{lazy'}$, at time t_k the events which triggered both those transitions will already be located in the input buffer (by the assumption above) and will trigger the transitions $q_i^{lazy'} \rightarrow q_{i+1}^{lazy'}$ and $q_{i+1}^{lazy'} \rightarrow q_{i+2}^{lazy'}$ respectively, removing themselves from the input buffer and adding themselves to the match buffer exactly as occurred during evaluation in A_{lazy}. Thus, at time t_l, both A_{lazy} and $A_{lazy'}$ contain s_u and s_v in their respective match buffers and do not contain s_u and s_v in their input buffers.

Thus, we have shown that in all cases the states of both A_{lazy} and $A_{lazy'}$ are identical at t_l. As stated above, this claim completes the proof.

Lemma 2 *Let P be a sequence pattern over event types $e_1, \cdots, e_k, \cdots, e_n$, with a single negated event e_k. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

We prove this lemma separately for two possible types of A_{lazy}, Post-Processing Chain NFAs and First-Chance Chain NFAs.

Let us first consider the case in which A_{lazy} is implemented as a Post-Processing Chain NFA. Let $S = s_1, ..., s_m$ be a stream of events. Assume without loss of generality that all events in S are within the time window W. If S does not contain an instance of a negated event e_k, then, by Lemma 1, it is accepted by A_{eager} if and only if it is accepted by A_{lazy}. Otherwise, we examine the two possible scenarios:

1. An event s_l of type e_k exists in S and satisfies the conditions required by the pattern. In this case, A_{eager} will discard all intermediate results and no match will be reported. In A_{lazy}, by construction, a state r exists, which is responsible for detecting e_k. If the execution for all matches never reaches r, then, since r precedes F, no instance will reach the accepting state and thus no match will be reported. Otherwise, if r is reached by some instance after s_l has arrived, then s_l will be located in the input buffer, by construction of a Post-Processing Chain NFA. Then, the edge e_r^{take} will take s_l and proceed to the rejecting state, thus discarding the instance. Finally, if r is reached by some instance before s_l has arrived, then e_k must appear at pattern end, after all positive events have already arrived. By construction, r will only have an outgoing edge $e_r^{timeout}$ to the accepting state, and the resulting *search_failed* event will not be processed. Following the arrival of s_l, the edge e_r^{take} will proceed to the rejecting state, since conditions for s_l are satisfied. Thus, in any case, A_{lazy} will discard all its partial matches.

2. An event s_l of type e_k exists in S and does not satisfy the conditions required by the pattern. In this case, A_{eager} will ignore s_l and report all matches. In A_{lazy}, by construction, the edge e_r^{take} of a state r will never be traversed, since its conditions do not hold for s_l. If r is reached by some instance after s_l has arrived, then s_l will be located in the input buffer, and a search will trigger the *search_failed* event, which will cause the instance to proceed to F. Otherwise, if r is reached by some instance before s_l has arrived, then e_k must appear at pattern end, after all positive events have already arrived. By construction, r will only have $e_r^{timeout}$ edge to the accepting state. When s_l arrives, the edge e_r^{take} will not be traversed, and the instance will only be affected by the *timeout* event, which will cause a transition to F. Thus, in any case, A_{lazy} will report all its positive matches.

To summarize, we have shown that in all cases A_{eager} and A_{lazy} report the same matches when applied on a stream S.

Now, let A_{lazy} be implemented as a First-Chance Chain NFA. Again, let $S = s_1, ..., s_m$ be a stream of events and assume w.l.o.g. that all events in S are within the time window W. If S does not contain an instance of a negated event e_k, then, by Lemma 1, it is accepted by A_{eager} if and only if it is accepted by A_{lazy}. Otherwise, we examine two possible scenarios:

1. An event s_l of type e_k exists in S and satisfies the conditions required by the pattern. In this case, A_{eager} will discard all intermediate results and no match will be reported. In A_{lazy}, by construction, a *take* edge e_k^{take} exists, detecting an event of type e_k and proceeding to the rejecting state R. Let q denote the source state of e_k^{take}. If the execution for all matches never reaches q, then, since q precedes F, no instance will reach the accepting state and no match will be reported. Otherwise, if q is reached by some instance after s_l has arrived, then s_l will be located in the input buffer, by construction of A_{lazy}. Then, e_k^{take} will take s_l and proceed to R. Note that q cannot be reached before s_l has arrived, since in this case e_k must appear at the pattern's end (by construction, requiring the positive event succeeding e_k to be taken prior to entering q), and this type of pattern is not supported by First-Chance Chain NFAs. Thus, in any case, A_{lazy} will discard all its partial matches.

2. An event s_l of type e_k exists in S and does not satisfy the conditions required by the pattern. In this case, A_{eager} will ignore s_l and report all matches. In A_{lazy}, by construction, e_k^{take} will never be traversed, since its conditions do not hold for s_l. With the exception of e_k^{take}, the structure of A_{lazy} is identical to that of a Lazy Sequence Chain NFA accepting a pattern P with e_k ignored. Therefore, by Lemma 1, all matches reported by A_{eager} will also be reported by A_{lazy}.

Thus, we have shown that in all cases A_{eager} and A_{lazy} report the same matches when applied on a stream S, which completes the proof.

Corollary 1 *Let P be a sequence pattern over event types $e_1, \cdots, e_n$, with an arbitrary number $1 \leq m < n$ of negated events. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

The proof is by induction on m, with Lemma 2 serving as an induction basis. The induction step for $m = i + 1$ is proven in the same way as Lemma 2, using the induction hypothesis instead of Lemma 1 for assuming the correctness of the pattern without the m^{th} negated event.

Lemma 3 *Let P be a sequence pattern over event types $e_1, \cdots, e_k, \cdots, e_n$, with a single iterated event e_k. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

We will prove this lemma by double inclusion, i.e., by showing that, given a potential match $M = s_1, ..., s_m$ for pattern P, A_{eager} accepts M if and only if A_{lazy} does. This, in turn, will be proven by induction on the number of events of type e_k in M.

For the induction basis, assume that M contains only a single event of type e_k. Then, let P' be a pattern identical to P, but with e_k as a non-iterated type. By Lemma 1, if M is a match for P, then it is also a match for P'. By iteration

definition, both A_{eager} and A_{lazy} also accept *P'*. Hence, either both NFA accept the match *M*, or both reject it.

For the induction step, assume that the condition holds for up to i events of type e_k, and let M contain *i+1* events of type e_k: $s_1^k, ..., s_i^k, s_{i+1}^k$.

Let M be accepted by A_{eager}. Let t_{i+1} be the arrival time of s_{i+1}^k. Then, a state q exists in Q_{eager}, such that at time t_{i+1} the NFA instance of A_{eager} is in q. Since $i \geq 1$, q is a state containing a self-loop taking e_k. Hence, before and after t_{i+1}, the current state was q. Since M is accepted, we know that all events that arrived after t_{i+1} caused the NFA instance to reach F from q. Let *M'* be a match identical to *M*, but without s_{i+1}^k. By the observation above, *M'* is also accepted by A_{eager}. By the induction hypothesis, *M'* is thus accepted by A_{lazy}. Now, since M is accepted by A_{eager}, it follows that s_{i+1}^k satisfies the conditions with other events in *M*. Therefore, when an *iterate* edge in A_{lazy} fetches a subset $s_1^k, ..., s_{i+1}^k$, it is added to the match buffer. By construction of a Lazy Iteration Chain NFA, the edge is traversed, reaching *F*. Hence, M is accepted by A_{lazy}.

Now, let M be accepted by A_{lazy}. Consequently, when a subset $s_1^k, ..., s_{i+1}^k$ is fetched by an iterate edge of A_{lazy}, all events in this subset satisfy the conditions with other events. In particular, it follows that a subset $s_1^k, ..., s_i^k$ satisfies the conditions as well. It also follows that during evaluation of M some NFA instance reaches the state with an outgoing *iterate* edge. Hence, during the same traversal which retrieves $s_1^k, ..., s_{i+1}^k$, a subset $s_1^k, ..., s_i^k$ is also generated and causes the transition to occur. That is, the match *M'* containing $s_1^k, ..., s_i^k$ without s_{i+1}^k is also accepted by A_{lazy}. By the induction hypothesis, *M'* is thus accepted by A_{eager}. By definition of the eager NFA, after s_i^k arrives during evaluation, an instance exists whose current state is a state q with a self-loop e_{loop} taking e_k. When s_{i+1}^k arrives, this instance will thus attempt a traversal of e_{loop}. By assumption, $s_1^k, ..., s_{i+1}^k$ satisfies all the conditions required by the pattern, hence this transition will succeed. Since *M'* is accepted, we know that all events that arrived after s_i^k in *M'* caused the NFA instance to reach F from q. Consequently, the same event sequence will cause the same transitions in *M*, i.e., M is accepted by A_{eager}.

To summarize, we have shown that, given a sequence pattern M containing an iterated event, A_{eager} accepts M if A_{lazy} does, and vice versa, which proves their equivalence.

Corollary 2 *Let P be a sequence pattern over event types $e_1, \cdots, e_n$, with an arbitrary number $1 \leq m < n$ of iterated events. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

The proof is identical to that of Corollary 1, using Lemma 3 instead of Lemma 2.

Definition 1 *Let $A = (Q, E, q_1, F, R)$ be an NFA. We will call A a **composite NFA of size K**, if there exist automata $A_1, \cdots A_k$, such that A consists of the union of them, with their respective initial states merged into q_1, the accepting states merged into F, and the rejecting states merged into R.*

Lemma 4 *Let A_1,A_2 be two composite NFAs of size K, i.e., A_1 consists of $A_1^1, \cdots A_k^1$ and A_2 consists of $A_1^2, \cdots A_k^2$. Then, if the sub-automata of A_1 and A_2 are pairwise equivalent, i.e., $\left(A_1^1 \equiv A_1^2\right) \wedge \cdots \wedge \left(A_k^1 \equiv A_k^2\right)$, then A_1 and A_2 are equivalent.*

The proof is by induction on K. For $K = 1$, the correctness follows immediately, since $A_1 = A_1^1 \equiv A_1^2 = A_2$. For the induction step, assume the claim to hold for $K = i$ and let A_1,A_2 be two composite NFAs of size $i + 1$. Now, assume that all sub-automata of A_1 and A_2 are equivalent. We will show the equivalence of A_1 and A_2 by mutual inclusion.

Let $M = s_1, ..., s_m$ be a potential match for A_1,A_2. Examine the following possible situations:

1. M is accepted by A_{i+1}^1. Then, by construction of A_1, it also accepts M. On the other side, by assumption of equivalence of A_{i+1}^1 and A_{i+1}^2, M is also accepted by A_{i+1}^2. By construction of A_2, it also accepts M.

2. M is accepted by A_1, but not accepted by A_{i+1}^1. Let $\bar{A}_1$,$\bar{A}_2$ be constructed by removing all the internal states of A_{i+1}^2 and A_{i+1}^2, respectively. M is accepted by $\bar{A}_1$, according to the assumption. Then, by the induction hypothesis, M is also accepted by $\bar{A}_2$, since $\bar{A}_1$,$\bar{A}_2$ are composite NFAs of size i. By construction of $\bar{A}_2$, M is also accepted by A_2.

3. M is not accepted by A_1. Then, in particular, it is not accepted by any of its sub-automata $A_1^1, \cdots A_i^1, A_{i+1}^1$. By assumption of equivalence, M is also not accepted by any of $A_1^2, \cdots A_i^2, A_{i+1}^2$. By construction of A_2, it also does not accept M.

To summarize, we have shown that, for any potential match M, it is either accepted by both A_1 and A_2 or by neither of them, which completes the proof.

Lemma 5 *Let P be a partial sequence pattern over event types $e_1, \cdots, e_n$, with possible negated or iterated events. Let A_{chain} be a Lazy Chain NFA for P, as presented in Section 3.3. Let $SEQ = \{seq_1, \cdots, seq_k\}$ be a set of all sequence patterns whose orders are allowed by P (i.e., P is a union of patterns in SEQ). Finally, let $A_{multi-chain}$ be a Lazy Multi-Chain NFA (Section 4.1), with each sub-chain being a Lazy Chain NFA for a sequence pattern in SEQ (Section 3.1). Then A_{chain} and $A_{multi-chain}$ are equivalent.*

The proof is by mutual inclusion. Let $M = s_1, ..., s_m$ be a potential match for A_1,A_2. Examine the following scenarios:

1. M is accepted by $A_{multi-chain}$. Then, in particular, it is accepted by at least one of its sub-chains. By construction, this Chain NFA A_i and A_{chain} are identical except for A_i having stricter ordering constraints. Consequently, M is also accepted by A_{chain}.

2. M is rejected by $A_{multi-chain}$. Then, it is also rejected by each of its sub-chains. Observe that the only difference between A_{chain} and each of

the sub-chains is the ordering constraint definition. Hence, the reason for the rejection can only be related to violating temporal conditions. Assume w.l.o.g. that an event s_p of type e_u in M precedes an event s_r of type e_v, which is forbidden by all chains. Also, assume that e_v precedes e_u in the frequency order (the opposite case is symmetrical). Then, as A_{chain} reaches a state responsible for taking e_u, s_r is already located in the match buffer. By definition of an ordering filter $succ_i$ for Lazy Chain NFA for partial sequences, $succ_i$ will contain e_v, hence s_p will not be taken by A_{chain}, discarding M.

Corollary 3 *Let P be a conjunction pattern over event types $e_1, \cdots, e_n$, with possible negated or iterated events. Let A_{chain} be a Lazy Chain NFA for P. Let us define $SEQ = \{seq_1, \cdots, seq_k\}$ as a set of all sequences over $e_1, \cdots, e_n$. Let $A_{multi-chain}$ be a Lazy Multi-Chain NFA with each sub-chain being a Lazy Chain NFA for a sequence pattern in SEQ. Then A_{chain} and $A_{multi-chain}$ are equivalent.*

The proof immediately follows from Lemma 5, since a conjunction pattern is an edge case of a partial sequence, with no ordering constraints between the primitive events.

Corollary 4 *Let P be a partial sequence pattern over $e_1, \cdots, e_n$. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

Assume w.l.o.g. that A_{eager} is implemented as a composite NFA, containing a sub-automaton for each valid sequence in P. An NFA with this structure necessarily exists, since each partial sequence can be represented as a disjunction of sequences. Let $A_{multi-chain}$ be a Lazy Multi-Chain NFA with each sub-chain being a Lazy Chain NFA for each valid sequence in P. By Lemma 4, A_{eager} and $A_{multi-chain}$ are equivalent. By Lemma 5, $A_{multi-chain}$ and A_{lazy} are equivalent. Then, by transitivity, A_{eager} and A_{lazy} are equivalent.

Corollary 5 *Let P be a conjunction pattern over event types $e_1, \cdots, e_n$. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

The proof immediately follows from Corollary 4.

Lemma 6 *Let P be a disjunction pattern over $e_1, \cdots, e_n$, i.e., $P = OR(p_1, \cdots p_m)$, where each p_i is a pattern over sequence, conjunction, partial sequence, negation and iteration operators. Let A_{eager} be an eager NFA and let A_{lazy} be a Lazy Multi-Chain NFA derived from P. Then A_{eager} and A_{lazy} are equivalent.*

Assume w.l.o.g. that A_{eager} is implemented as a composite NFA, containing a sub-automaton for each sub-pattern p_i. Then, A_{eager} and A_{lazy} are both composite NFA of size m. By definition of p_i, by Lemmas 1,2 and 3 and by

Corollaries 4 and 5, each sub-automaton for p_i in A_{eager} is equivalent to a corresponding sub-chain in A_{lazy}. Then, by Lemma 4, A_{eager} and A_{lazy} are equivalent.

We will now complete the proof of Theorem 1. If P does not contain the disjunction operator, the equivalence of A_{eager} and A_{lazy} follows from Lemmas 1,2 and 3 and from Corollaries 4 and 5. Otherwise, let P' be the DNF form of pattern P. It is sufficient to prove that the claim holds for P'. By definition of DNF, P' is a disjunction pattern. Thus, by Lemma 6, the corresponding eager and lazy NFA for P' are equivalent.■

6 Experimental Evaluation

This section presents a detailed experimental analysis of our method. We applied Lazy (Multi-)Chain NFA on wide range of patterns, assessing their efficiency and scalability. A comparison with an eager NFA-based CEP framework was conducted, demonstrating an overall performance gain of up to two orders of magnitude in most scenarios.

Our metrics for this study are throughput, memory consumption, and runtime complexity. Throughput is defined as the number of events processed per second. For memory consumption, we estimate the peak size (in MB) of the reserved system memory. Runtime complexity is measured as the number of calculations and memory modifications per event. For convenience, these two types of operations were evaluated separately.

Real-world historical stock price data from the NASDAQ stock market, taken from [1], was used in the experiments. This data spans a 1-year period, covering over 2100 stock identifiers with prices updated on a per minute basis. Our input stream contained 80,509,033 primitive events, each consisting of a stock identifier, a timestamp, and a current price. We augmented the event format with additional attributes, namely a history list of older prices and an identifier of a region to which the stock belongs, e.g., North America or Europe. For each of the 8 regions defined by NASDAQ, a corresponding event type was defined. A primitive event was then considered to belong to a type corresponding to the region of its stock identifier.

The patterns used during evaluation differed only in their operator (e.g., SEQ or AND), the participating event types (i.e., regions) and the time window size. All patterns shared an identical condition of high correlation between each pair of primitive events. That is, the Pearson correlation coefficient between price history lists was required to exceed a predefined threshold T. For example, for a sequence operator, three event types, *NACompany*, *EuCompany*, *AfrCompany*, and a time window h, the corresponding pattern would be declared as follows:

$$
\begin{aligned}
&PATTERN\ SEQ \\
&\quad (NACompany\ a, EUCompany\ b, AfrCompany\ c) \\
&WHERE\ skip_till_any_match\ \{ \\
&\quad Corr\,(a.history, b.history) > T \\
&\quad AND\ Corr\,(b.history, c.history) > T \\
&\quad AND\ Corr\,(c.history, a.history) > T\} \\
&WITHIN\ h.
\end{aligned}
\tag{10}
$$

The following pattern types were used during this study:

- Sequences of 3, 4 and 5 primitive events (marked on all graphs as SEQ3, SEQ4 and SEQ5, respectively).
- Conjunctions of 2 and 3 primitive events (marked as AND2, AND3). Due to the extremely rapid growth of NFA instances during eager evaluation attempts, experiments for higher numbers of events could not be conducted in the environment available for this research.
- Negation patterns for SEQ3, AND3 and SEQ5. In all of the above, the second event was negated. For SEQ5 alone, we evaluated both methods for lazy negation presented in Section 3.4. They are marked as Lazy-PP and Lazy-FC on all graphs. For the rest of the patterns, only Lazy Post-Processing Chain NFA was used.
- Iteration pattern of 3 events (i.e., SEQ(a,b+,c)), with the second event being iterated (marked as ITER3). Due to the exponential number of instances and limited hardware resources, experiments for this pattern were conducted on time windows half the size of those used for other pattern types.
- Disjunction patterns consisting of: (1)two sequences of two events (OR2SEQ2); (2)four sequences of two events (OR4SEQ2); (3)two sequences of four events (OR2SEQ4).

All NFA models under examination were implemented in Java. The experiments were run on a machine with 2.20 Ghz CPU and 16.0 GB RAM.

6.1 Time Window Size

In our first set of experiments, we compared how the lazy and the eager strategies scale with a growing time window size. Both algorithms were repeatedly applied on a data stream, with values for h ranging between 5 and 30 minutes. For the iteration pattern alone, the range was modified to 2-12 minutes.

All figures summarizing this experiment contain several sub-graphs, presenting results for different pattern types as listed above.

Figure 8 describes the throughput as a function of h. For presentation clarity, the logarithmic scale was used. A steady increase of one to two orders of magnitude can be observed for the lazy approach as compared to its eager counterpart.

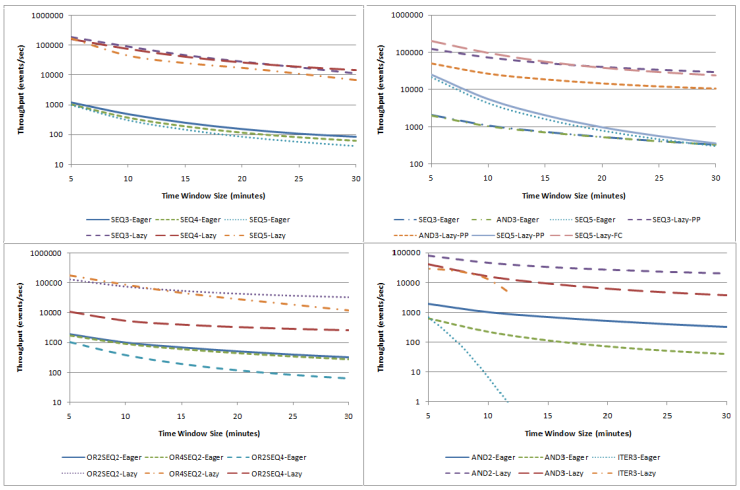

Figure 8: Throughput as a function of time window size for various patterns (logarithmic scale).

The performance decline for larger values of h is generally smoother for lazy NFA or comparable in both models. For the negated sequence of five events, a clear advantage of the First-Chance Chain NFA over the Post-Processing method can be seen. This holds due to the negated event being the second in the sequence. Hence, according to First-Chance strategy, verification of its absence can occur as the third event is accepted, rather than after all positive events have arrived. For disjunctions, sub-sequence length affects the overall throughput significantly more than the number of sub-sequences.

Memory consumption measurements are presented in Figures 9 and 10. Figure 9 compares the total memory used during eager and lazy evaluation of various patterns. For all pattern types, lazy automata require less memory than eager automata. This is most evident for sequences and disjunctions, for which the lazy solutions are 1.5 to 3 times more economical for larger time windows. The smallest gain is observed for conjunctions, which can be explained by the absence of ordering filters for this pattern type (Section 3.2).

It should be noted that the above figure only displays the peak memory usage, i.e., the maximal recorded amount of required memory. However, the difference in average memory consumption is more critical. Whereas an eager NFA maintains the same high amount of reserved memory most of the time, a Lazy Chain NFA only achieves its peak during brief 'bursts'. These bursts follow an arrival of a rare event, which triggers a search in the input buffer and creates a large number of intermediate instances.

Two main types of objects contribute to the overall memory consumption. First, new NFA instances are created for each partial match. Second, in lazy NFA primitive events are buffered upon arrival. While the lazy evaluation method keeps more buffered events, it creates dramatically fewer instances.

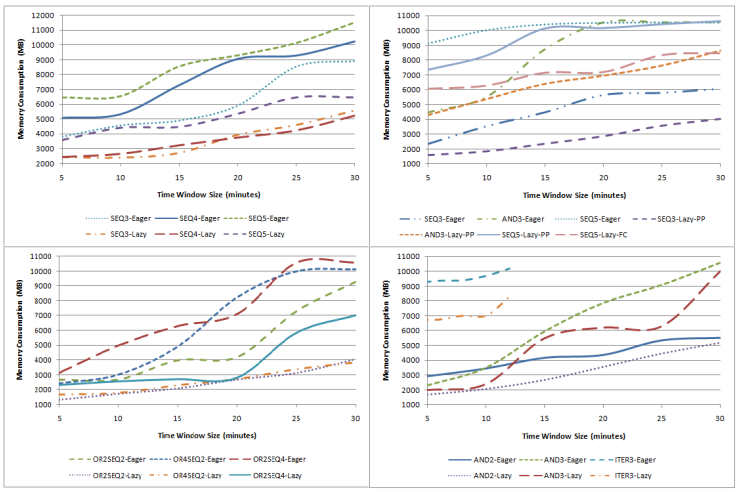

Figure 9: Peak memory consumption as a function of time window size.

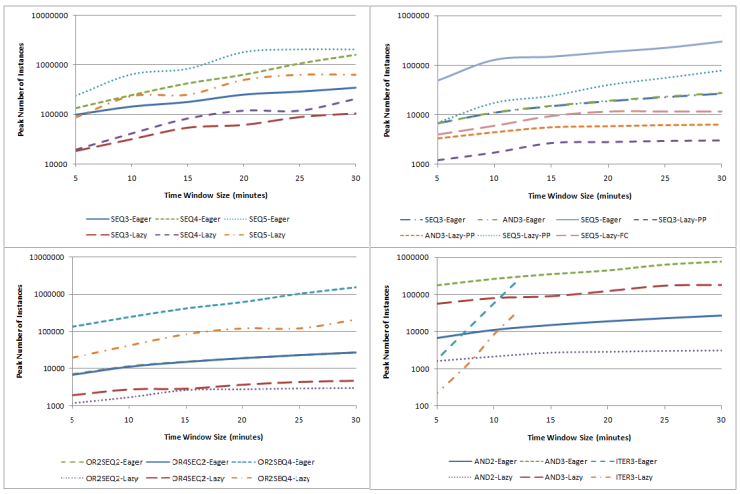

Figure 10: Peak number of NFA instances as a function of time window size (logarithmic scale).

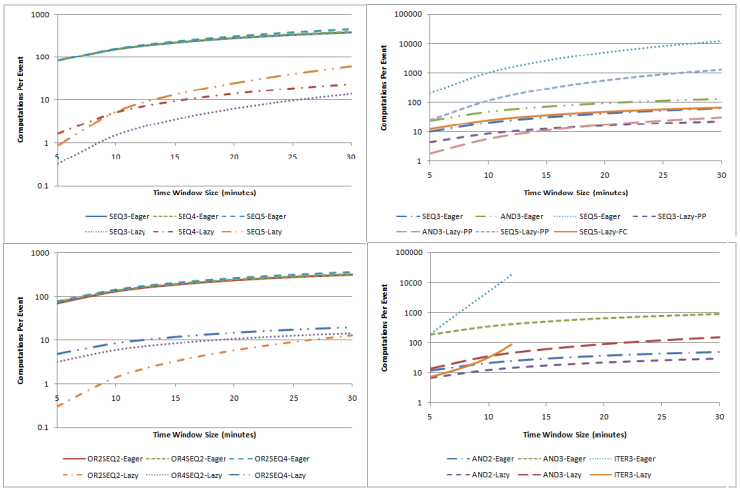

Figure 11: Number of computations per event as a function of time window size (logarithmic scale).

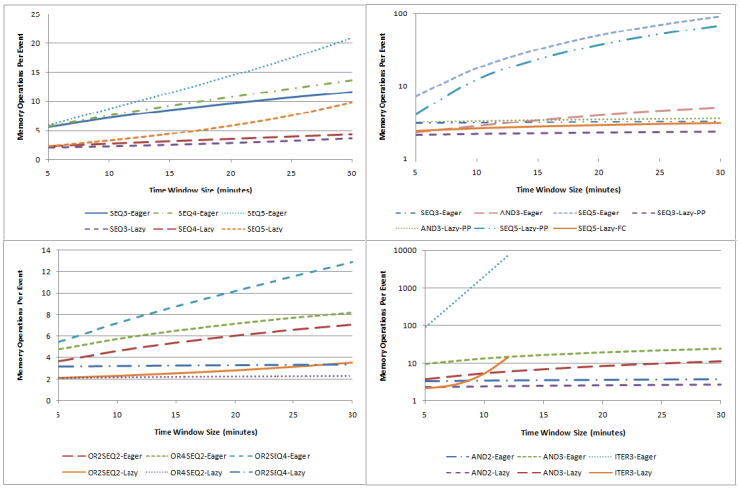

Figure 12: Number of memory operations per event as a function of time window size (logarithmic scale for negations, conjunctions and iterations).

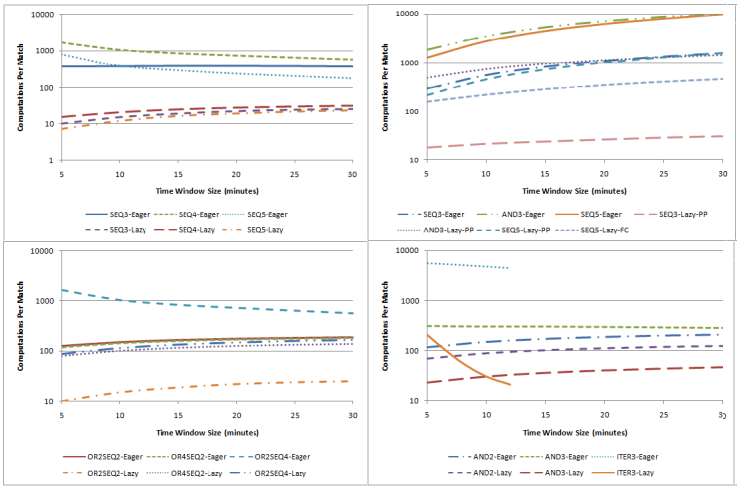

Figure 13: Number of computations per pattern match as a function of time window size (logarithmic scale).

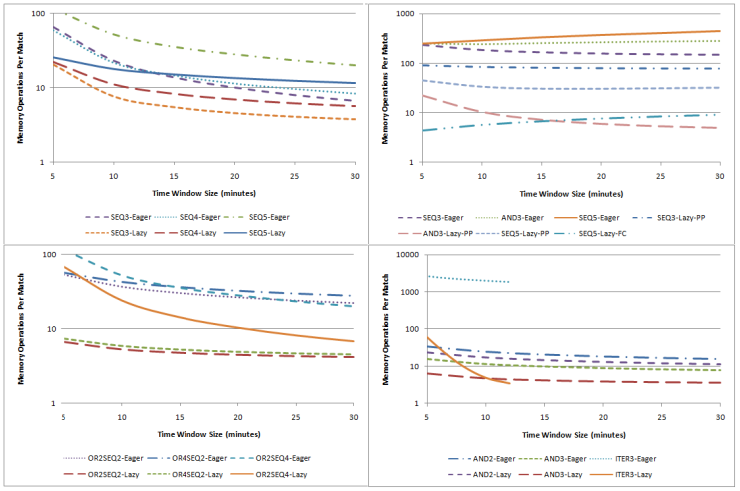

Figure 14: Number of memory operations per pattern match as a function of time window size (logarithmic scale).

Figure 10 demonstrates this by comparing the peak number of simultaneously active instances during evaluation. Again, the logarithmic scale was used. On average, a Lazy (Multi-)Chain NFA maintains 5 to 100 times fewer instances at the same time. The biggest gain was achieved for large negation and disjunction patterns.

Now we proceed to the runtime complexity comparison. In this experiment, we analyzed the two evaluation mechanisms in terms of operations performed per primitive event and per successful pattern match. These operations can be divided into two separate categories. The first includes the calculations executed during predicate evaluation. Memory modifications, such as NFA instance creation and deletion, buffering primitive events and removing them upon evaluation, comprise the second category.

In Figure 11 the number of predicate evaluations for each tested pattern can be seen. On average, lazy NFAs execute 10 to 200 times fewer computations per primitive event. This difference tends to be larger for more complex patterns involving more event types.

Figure 12 presents the results for memory modifications. Despite the need to maintain a complex data structure for input buffering and to perform a large volume of insertions, searches and removals, the lazy method does not suffer any performance penalty in comparison to the eager approach. Furthermore, for most patterns, significantly fewer operations were recorded during lazy evaluation. This can be explained by the drastically reduced number of instances to be created (and hence destroyed).

Finally, Figures 13 and 14 display the runtime complexity measurements adjusted per detected pattern match. The results are similar to those demonstrated above for per event operations.

6.2 Rarest-to-most-frequent Event Ratio

The core principle of the lazy evaluation mechanism is to re-order the input stream so that less frequent events will be processed earlier. Thus, from an analytical standpoint, this method achieves its biggest performance gain when the frequencies of the participating events are highly varied. As a consequence, the impact of this re-ordering will diminish as events arrive at at more similar rates, and will completely vanish if all event types are of identical frequency. An interesting question is, how does a Lazy Chain NFA perform in these scenarios as compared to an eager NFA?

In our next experiment we attempted to answer this question. To this end, we evaluated a subset of patterns presented above on different selections of event types, using a parameter called *rarest-to-most-frequent event ratio*. It is defined as the maximal ratio of arrival rates of two event types in a pattern.

The frequency of each type was estimated according to the overall number of companies belonging to the respective region. For example, while NASDAQ has 267 registered companies located in Europe, there are only 8 such companies in Africa. Hence, the ratio of the event type *AfrCompany* to *EuCompany* is 8:267, or approximately 1:33. By carefully choosing sets of event types, we created a

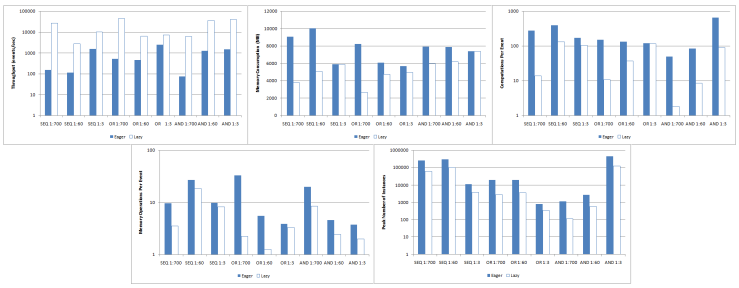

Figure 15: Eager vs. Lazy evaluation performance for patterns with varied rarest-to-most-common event ratio: (a)throughput; (b)peak memory consumption; (c)number of predicate evaluations per event; (d)number of memory modifications per event; (e)peak number of simultaneously active instances.

number of patterns with a rarest-to-most-frequent event ratio ranging from 1:3 to 1:700. Each of the patterns was then evaluated with time window size set to 20. All performance metrics presented in the beginning of this section were measured during this experiment.

The results are displayed in Figure 15. Note that in this case we are not interested in the absolute values. Instead, we examine the performance difference of the eager and the lazy evaluation mechanisms. As expected, this difference is inversely proportional to the rarest-to-most-frequent ratio. For small values of this ratio, the results for the lazy method are always at most equal to those of the eager method.

6.3 State-of-the-art Comparison

In our last experiment we compared the performance of the Lazy Chain NFA with the official implementation of SASE+[30]. The set-up described in Section 6.1 was used, in which we studied the performance of both systems as a function of time window size. Sequence, negation and iteration patterns were employed. To avoid modifying the code of SASE+, we only measured throughput during the evaluation.

Figure 16 shows the results. Since our implementation of the eager NFA was based on SASE+ and carefully followed its formal description, a high degree of similarity with Figure 8 can be observed. As demonstrated earlier, using the Lazy Evaluation method yields a performance gain of one to two orders of magnitude in terms of processed events per time unit.

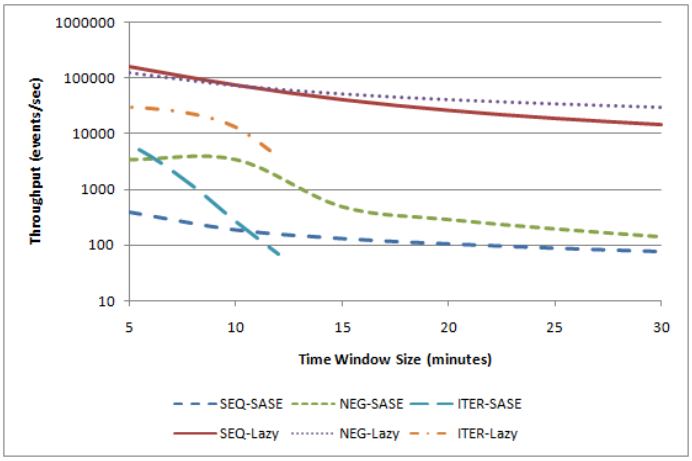

Figure 16: Throughput of SASE+ and the Lazy Chain NFA for sequence, negation and iteration patterns (logarithmic scale).

7 Related Work

The detection of complex events has become a very active research field in recent years [14]. Modern CEP frameworks developed as a direct descendant of earlier Data Stream Management Systems [6, 11, 19], which focus on processing data coming from continuous, usually multiple input streams, whereas the goal of CEP is to detect events: occurrences of complex situations of interest. Apart from SASE [28] and SASE+ [3], thoroughly discussed above, numerous event processing mechanisms were developed by different research communities [2, 8, 13, 18, 24].

CEP systems are generally based on SQL-like specification languages. CEL [15, 16] is a declarative language used by the Cauyga system. It supports patterns with Kleene closure and event selection strategies. CQL [5] made it possible to create transformation rules with a unified syntax for processing both information flows and stored relations. TESLA [12] combines high expressiveness with a relatively small set of operators. It offers completely programmable per-event selection policies. CEDR [7] presents a complex and flexible temporal model, based on three different timings. Our lazy evaluation framework is based on SASE+ [20], an expressive language, including constructs such as iterations and aggregations. Nevertheless, the operators it supports serve as basic blocks for most of the described languages.

Optimizing event processing performance by pattern re-writing (re-ordering) techniques has been discussed in numerous studies. [25] and [26] employ formally defined cost models to derive the best possible evaluation plan, according to which input patterns are re-organized. The objectives of query re-writing are throughput and latency balancing. Both frameworks consider distributed CEP environment.

The concept of lazy evaluation has been widely applied in multiple research fields. Related work was also conducted in the context of complex event processing. ZStream [23] uses tree-based query plans to represent complex patterns. As the most efficient plan is chosen, the pattern detection order is dynamically adjusted by internal buffering. In [4], the authors describe "plan-based evaluation", where temporal properties of primitive events are exploited to reduce network communication costs. [10] presents an XPath-based mechanism for filtering XML documents in stream environments. This method postpones costly operations as long as possible. While many of the ideas discussed by the aforementioned studies are close to ours, none considers the finite automata environment, which is the primary focus of our work.

Several authors propose postponing and pre-processing mechanisms for NFA-based evaluation frameworks. In [31], the authors present an optimization method based on sharing common operations between instances and postponing per-instance part. The focus of that technique is on improving the performance of Kleene closure patterns under the skip-till-any-match selection policy. [17] discusses a mechanism similar to ours, including the concept of storing incoming events and evaluating more selective events before more frequent ones. However, the authors do not consider complex pattern types, such as negation and iteration. [9] proposes a matching strategy including a pre-processing step, which is similar to our input buffering. It is used for applying pruning techniques, such as filtering and partitioning, rather than for event reordering. APAM [29] is a hybrid eager-lazy evaluation method. For each event type, it determines which method to use for minimizing detection latency.

8 Conclusions

In this paper we have presented a generic NFA topology, denoted as Lazy Chain NFA, which implements the concepts of lazy evaluation and event reordering for efficient detection of complex patterns of various types. It employs internal event buffering to process events in ascending order of their frequencies, instead of using the order imposed by the input pattern. Extensive experimental evaluation results demonstrated significant improvement of this solution over the eager evaluation mechanism in terms of throughput, memory consumption, and runtime complexity. Our future work will explore several research avenues, such as efficient multi-query support, handling uncertainty, and adapting our techniques to a distributed environment.

9 Acknowledgments

The research leading to these results has received partial funding from the [European Union's] Seventh Framework Programme [FP7-ICT-2013-11] under grant agreements n. [619491] (FERARI), n. [619435] (SPEEDD) and n. [688380] (VAVEL).

References

[1] http://www.eoddata.com.

[2] A. Adi and O. Etzion. Amit - the situation manager. *The VLDB Journal*, 13(2):177–203, 2004.

[3] J. Agrawal, Y. Diao, D. Gyllstrom, and N. Immerman. Efficient pattern matching over event streams. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, pages 147–160.

[4] M. Akdere, U. Çetintemel, and N. Tatbul. Plan-based complex event detection across distributed sources. *Proc. VLDB Endow.*, 1(1):66–77, 2008.

[5] A. Arasu, S. Babu, and J. Widom. The cql continuous query language: Semantic foundations and query execution. *The VLDB Journal*, 15(2):121–142, 2006.

[6] H. Balakrishnan, M. Balazinska, D. Carney, U. Çetintemel, M. Cherniack, C. Convey, E. Galvez, J. Salz, M. Stonebraker, N. Tatbul, R. Tibbetts, and S. Zdonik. Retrospective on aurora. *The VLDB Journal*, 13(4):370–383, 2004.

[7] R. S. Barga, J. Goldstein, M. H. Ali, and M. Hong. Consistent streaming through time: A vision for event stream processing. In *CIDR*, pages 363–374, 2007.

[8] L. Brenna, A. Demers, J. Gehrke, M. Hong, J. Ossher, B. Panda, M. Riedewald, M. Thatte, and W. White. Cayuga: A high-performance event processing engine. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, pages 1100–1102. ACM.

[9] B. Cadonna, J. Gamper, and M. H. Böhlen. Efficient event pattern matching with match windows. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 471–479. ACM, 2012.

[10] C. Y. Chan, P. Felber, M. N. Garofalakis, and R. Rastogi. Efficient filtering of xml documents with xpath expressions. *VLDB J.*, 11(4):354–379, 2002.

[11] J. Chen, D. J. DeWitt, F. Tian, and Y. Wang. Niagaracq: A scalable continuous query system for internet databases. *SIGMOD Rec.*, 29(2):379–390, 2000.

[12] G. Cugola and A. Margara. Tesla: a formally defined event specification language. In *DEBS*, pages 50–61. ACM, 2010.

[13] G. Cugola and A. Margara. Complex event processing with t-rex. *J. Syst. Softw.*, 85(8):1709–1728, 2012.

[14] G. Cugola and A. Margara. Processing flows of information: From data stream to complex event processing. *ACM Comput. Surv.*, 44(3):15:1–15:62, 2012.

[15] A. Demers, J. Gehrke, M. Hong, M. Riedewald, and W. White. Towards expressive publish/subscribe systems. In *Proceedings of the 10th International Conference on Advances in Database Technology*, pages 627–644. Springer-Verlag.

[16] A. Demers, J. Gehrke, and B. Panda. Cayuga: A general purpose event monitoring system. In *CIDR*, pages 412–422, 2007.

[17] C. Dousson and P. Le Maigat. Chronicle recognition improvement using temporal focusing and hierarchization. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 324–329, 2007.

[18] B. Gedik, H. Andrade, K. L. Wu, P. S. Yu, and M. Doo. Spade: the system s declarative stream processing engine. In *SIGMOD Conference*, pages 1123–1134. ACM, 2008.

[19] The STREAM Group. Stream: The stanford stream data manager. Technical Report 2003-21, Stanford InfoLab, 2003.

[20] D. Gyllstrom, J. Agrawal, Y. Diao, and N. Immerman. On supporting kleene closure over event streams. In *ICDE*, pages 1391–1393. IEEE, 2008.

[21] L. Harada and Y. Hotta. Order checking in a cpoe using event analyzer. In *Proceedings of the 14th ACM International Conference on Information and Knowledge Management*, pages 549–555. ACM, 2005.

[22] I. Kolchinsky, I. Sharfman, and A. Schuster. Lazy evaluation methods for detecting complex events. In *Proceedings of the 9th ACM International Conference on Distributed Event-Based Systems*, pages 34–45. ACM, 2015.

[23] Y. Mei and S. Madden. Zstream: a cost-based query processor for adaptively detecting composite events. In *Proceedings of the 29th ACM SIGMOD Conference*, pages 193–206. ACM, 2009.

[24] P. R. Pietzuch, B. Shand, and J. Bacon. A framework for event composition in distributed systems. In *Proceedings of the ACM/IFIP/USENIX 2003 International Conference on Middleware*, pages 62–82. Springer-Verlag, 2003.

[25] E. Rabinovich, O. Etzion, and A. Gal. Pattern rewriting framework for event processing optimization. In *Proceedings of the 5th ACM International Conference on Distributed Event-based System*, pages 101–112. ACM, 2011.

[26] N. P. Schultz-Møller, M. M., and P. R. Pietzuch. Distributed complex event processing with query rewriting. In *DEBS*, 2009.

[27] F. Wang and P. Liu. Temporal management of rfid data. In *Proceedings of the 31st International Conference on Very Large Data Bases*, VLDB '05, pages 1128–1139. VLDB Endowment, 2005.

[28] E. Wu, Y. Diao, and S. Rizvi. High-performance complex event processing over streams. In *SIGMOD Conference*, pages 407–418. ACM, 2006.

[29] I. Yi, J. G. Lee, and K. Y. Whang. Apam: Adaptive eager-lazy hybrid evaluation of event patterns for low latency. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 2275–2280. ACM, 2016.

[30] H. Zhang, Y. Diao, and N. Immerman. Recognizing patterns in streams with imprecise timestamps. *Proceedings of the VLDB Endowment*, 3(1-2):244–255, 2010.

[31] H. Zhang, Y. Diao, and N. Immerman. On complexity and optimization of expensive queries in complex event processing. In *SIGMOD*, pages 217–228, 2014.